LIBRARY OF CONGRESS

in collaboration with

LAWRENCE BERKELEY NATIONAL LABORATORY

+ + + + +

AUDIO MEDIA PRESERVATION THROUGH IMAGING CONFERENCE

+ + + + +

THURSDAY
JULY 16, 2015

+ + + + +

The Conference met in the Mumford Room, Library of Congress, 101 Independence Avenue, S.E., Washington, D.C., at 9:30 a.m.

ORGANIZING COMMITTEE

PETER ALYEA, Library of Congress
EUGENE DEANNA, Library of Congress
CARL HABER, Lawrence Berkeley National Laboratory
ADRIJA HENLEY, Library of Congress
STEPHEN LEGGETT, Library of Congress

SPEAKERS

MARK SWEENEY, Associate Librarian, Library Services
OTTAR JOHNSEN, Haute ecole d'ingenieurs et d'architectes de Fribourg
STEFANO SERGIO CAVAGLIERI, Swiss National Sound Archives
JOHN MCBRIDE, University of Southampton
STIG L. MOLNERYD, National Library of Sweden
JOSHUA STERNFELD, The National Endowment for the Humanities

JESSE JOHNSTON, The National Endowment for the Humanities
BILL VEILLETTE, Northeast Document Conservation Center
MASON VANDER LUGT, Northeast Document Conservation Center
JAMES NYE, University of Chicago Library
SUNDAR GANESAN, Roja Muthiah Research Library
SURESH BABU GOVINDARAJU CHANDRAN, Roja Muthiah Research Library
CARLENE STEPHENS, National Museum of American History
SHARI STOUT, National Museum of American History
FENELLA FRANCE, Library of Congress
BILL KLINGER, Klinger Engineering Services

TABLE OF CONTENTS

P-R-O-C-E-E-D-I-N-G-S

(9:41 a.m.)

MR. ALYEA: So, welcome. This is the first meeting of this particular group of people. This is a very new technology. We're really excited that all of you have been able to come and we hope that we continue this dialog even past this conference. So we have some opening remarks from Mark Sweeney.

MR. SWEENEY: Thank you, Peter, and good morning to all. I'm Mark Sweeney, the Associate Librarian for Library Services, and on behalf of the Librarian of Congress, Dr. James Billington, it's my pleasure to welcome you today to the Library and to this three-day conference on Audio Media Preservation Through Imaging Technology.

As Peter says, this is just an excellent opportunity for us to share our knowledge and experience and sort of chart a path for the future. And we're having wonderful weather today here in Washington and for those that had an opportunity to go out to Culpeper, I think you saw something really special yesterday.

So audio recording has existed for over 150 years. The development from an phonautograph, to tin foil, to the myriad of today's, you know, modern high-fidelity digital recordings, you know, it's all amazing and astounding.

And over a comparatively short period of time, about 15 years, many of those in attendance here have developed, introduced, refined, the imaging technology to capture audio from physical audio carriers for the purpose of preserving information from these valuable artifacts.

The collaboration that has surrounded the development of these specialized imaging tools has been open and a collegial one, and there have been substantial contributions from the Lawrence Berkeley National Laboratory, the University of Applied Science in Fribourg, Switzerland, the University of California, Berkeley, Northeast Document Conservation Center, First Sounds, the National Museum of American History, the Thomas Edison National Historical Site, Roja Muthiah

Research Library in Chennai, India, the University of Chicago, and of course as well, the Library of Congress here today.

The research has received major funding from a number of sources, the Institute for Museum and Library Services, the National Endowment for the Humanities, the Library of Congress, as well as support from the University of California, the Department of Energy, the Smithsonian Institution, the Schenectady Museum of Innovation and Science, the National Archives and Records Administration, Harvard University, the Council on Library and Information Resources, the Mellon Foundation, Guggenheim Foundation, the MacArthur Foundation, so we have many supporters that have been helping us along the way.

As is true with any technological development, there are successes and challenges. Utilization imaging technology on audio media has already given us access to recordings that were broken into pieces and too fragile of a state to be touched physically, and in early experimental formats for which, you know, we have no playback equipment.

So only a relatively short time ago, these recordings were considered lost to history, but now they're accessible again. But there are also issues that have yet to be solved. I would put high on this list, maturation of this technology to meet our specific needs at scale.

This gathering of collection holders and instrument builders can help define the priorities for new tools to meet our needs. By balancing the most pressing needs of collections with the capabilities of the technology today, while also planning for the future, those gathered here can help shape the development and the roadmap for these imaging media systems of tomorrow.

There are so many sound recordings in need of preservation, setting the proper priorities for development is the first imperative. We believe this technology has diverse uses as a non-invasive format migration tool for media in a wide range of physical conditions. It also serves as an assessment and measurement tool to set standards for media playback. These are core requirements for preservation.

A future roadmap will allow others in the field to consider how this technology fits into their preservation strategies and to plan accordingly. As the technology gains capabilities and preservation specialists accumulate real-world experience with these tools, the position and acceptance of the technology will broaden; should broaden.

By working cooperatively, this group of mavericks has made considerable progress over the relatively short period of time, maturing the application of new technology for a very specific purpose. As future challenges are met and solutions are proposed, the collegial nature of this group should see it through the most difficult future issues.

Since preservation serves not only living scholars and patrons, but also future generations, the generosity of this group shows that its deep knowledge to solve long-standing problems inherent with this valuable historical record is recognized, so I thank you for your dedication.

I'd like to conclude by also thanking the members of our conference organizing team. That would be Peter Alyea, Gene DeAnna, Carl Haber, Adrija Henley, Steven Leggett, and Angela Newburn. Thank you for your hard work. Please enjoy the conference and I look forward to sitting in and learning more. Thank you.

MR. DEANNA: Good morning, everyone. I'm Gene DeAnna. I'm head of the recorded sound section here at the Library, and I want to welcome all of you that I didn't welcome yesterday at the tour of the Packard Campus. This is a great, exciting meeting. It's wonderful to see you all here today.

We had a really, I think, interesting tour, we had great discussions yesterday, or beginning discussions, at lunch, in the hallways between our stops at the Packard Campus, and before we get too far along, I want to make sure to invite all of you who weren't able to make the trip to please let us know and we'll try to make arrangements to get you out there to see the facility. I'm sure it'll be worth your while.

I want to second Mark's thanks, particularly to Adrija Henley and my colleague, Steve Leggett, who I'm not sure has made it today, but Steve is the coordinator of our National Film Board and our National Recording Preservation Board, and there are several board members here today.

The support of the Library and Congress of these boards has allowed them to support national efforts at preserving film and sound. You know, it helps to make events like this possible. As an archivist who's been in the field for over 25 years now, seeing this gathering of scientific, technical, and archival expertise focused on an exciting and still emerging technology for capturing sound is extremely gratifying.

In December of 2012, the Library and the National Recording Preservation Board, with the assistance of the Council of Library Information Resources, published a national plan for the preservation of our recorded legacy. That plan distilled over 100 preliminary recommendations gathered from an extensive study, multi-year study, and the recommendations of an eminent group of expert task forces.

Distilled them down to about 32 major recommendations and since its publication, that plan has been anything but just another report of good idea that sits on a shelf somewhere. It's been a vital document. It has influenced the direction of the work and provided a national framework for the focus of audio preservationists and the institutions and organizations that support their work all across this country, and actually, internationally as well.

So today, we gather here under the auspices of Recommendation 1.7 of the National Recording Preservation Plan, which says to encourage scientific and technical research leading to the development of new technologies to recover, reformat, and preserve audio recording media.

The development of imaging technology for audio is a terrific story. When you tell people about it and they read about, they respond to it like it's magic, and there is an element of magic to it, I think. And, you know, from a more technical standpoint, it's just really cool. There remains a lot of work to be done, discussing, defining, guiding that work.

Defining the framework of how to guide that work is going to be one of the important goals of this conference. Now, for me as an archive manager of a too large archive of recorded sound, I see three clear significant benefits and pathways to development to audio preservation that imaging technology offers, and that does not include the coolness factor.

The first is obviously the safe recovery of audio from fragile, damaged, or broken carriers, sound that otherwise might not be recovered, so playback of the unplayable or probably should not be played. And that also includes the safe playback of unclean, unprepped recordings to enable us to identify them for purposes of assessment and prioritization.

You'd be amazed at the resources that have to go into the conservation prep work to get recordings that, you just want to hear what they are, and get them ready to put a stylus on them, it's commitment of a lot of resources and you can imagine those materials that demand that kind of work often get just shunted aside for practical purposes.

Secondly, it's a non-time-based method for making transfer of groove media and it can be carried out in parallel and an unmonitored workflow. So I decided to do a little third-grade math, because occasionally, I like to push beyond my comfort zone, and I calculated that LC holds about, oh, say, a million hours of unique audio, that is, no dupes. One-of-a-kind audio. So clearly, if you think about a million hours of preservation to do in standard time-based capture, we need to develop other ways of doing this that overcome playback time and the limitations of standard reformatting.

And thirdly, this is the little harder for us, this is new for me to get my head around this, but it's the transformational ability that 3-D imaging provides for digitally preserving the carriers themselves, that is, the grooves, to image those grooves, capture that image for future generations. Audio archivists operate, and justifiably so, under the understanding that all of this media is going to degrade or eventually become unplayable for one reason or another, obsolescence, over time. But imaging is adding that new paradigm.

The goal of capturing and preserving the groove structure of historical recordings, historically important recordings, the earliest recordings of humanity, is really now a realistic consideration for archivists, and I guess we have to ask ourselves what's the worth, and what's that going to be worth down the line in a couple hundred years for people?

So there's still a lot of work to do, but one thing I want to point out, yesterday at the tour, the group saw what I think is one of the great success stories of A/V preservation. Up on the third floor in the video lab there's an array of five SAMA robots, they work in parallel, their hopper's full of pneumatic, VHS, and beta cartridge tapes.

These are video formats that carry the vast majority of recorded television worldwide. So 20-some decks digitizing in parallel in a workflow that generate now 25 terabytes of content a week. Ladies and gentlemen, I can tell you that ten years ago, television preservation was non-existent, virtually non-existent.

This content was more at risk than our nitrate, more at risk than our audiotape, our lacquered disks, and our wax cylinders. Now, my colleagues in the moving image section and one yesterday, while he talked to the group, mentioned the word, finishing, and that's a word that doesn't come up too often in the recorded sound section.

Finishing preserving television, and you think, well, what happened? Well, obviously, the Packard Humanities Institute and the Library of Congress, and Congressional support happened, but really what happened to make that a reality was good science, engineering brilliance, and a clear understanding of the archival principles and practices of preservation.

So scientists, engineers, and archivists, which is exactly the mix we see in this room today. So I look forward to this conference. I think it's very, very important, essentially, that we continue down the development of imaging and I want to welcome you, and I look forward to the discussions. So now I want to introduce my colleague, Peter Alyea, who leads the IRENE project here at the Library of Congress. Peter.

MR. ALYEA: Good morning. I don't have a lot to say, but I wanted to give a little perspective that I think may not be the common way we think about this. There are, obviously, a lot of issues involved in this preservation, there's degradation of the materials, and there's the time factor of getting all these things preserved.

But this imaging technology actually does an interesting thing. When I came to the Library of Congress as a preservation specialist preserving audio recordings, we went to analog tape, we preserved things, and actually, that's incorrect to say. We did not actually, technically, preserve things. At that time, preserving something was making sure the original carrier was in good condition. It was very much from the book world.

What I did was reformat materials and it wasn't considered part of preservation. Not so long ago, people realized as these materials are degrading and going away, and equipment to transfer them was becoming less accessible, that the process of reformatting was actually part of preservation, because you might lose the original item.

However, that did, for the particular items we're talking about today, that actually does a transformation. These items are physical items and when you do a transfer, you actually convert this physical thing into an audio stream and then you recapture that audio stream. So the thing that's so fascinating about the idea of imaging is we're back to the idea of actually preserving the original item again, and it gave us a whole host of possible uses in the future for these kinds of things.

I thought I'd mention a couple technologies that have been developed over a period of time, which aren't actually in the physical realm, but kind of highlight this. There's a technology called, from plangent processes, which uses bias tones, which are high-frequency signals on audiotapes, which aren't for listening back, but for use for alignment of the tape, to actually re-speed the tape.

You can actually lock-on to these signals and then actually get a better speed for the tape playback. That's an idea where you're using something that's not normally preserved on a tape for a purpose of actually preserving the thing better. There's also work at the University of

Maryland where they're actually marrying records of the power grid cycles to signals on audio recordings because there is date stamps on the actual fluctuation in power grids and you can actually then match these up with audio recordings and figure out exactly when something was recording it.

So those may seem very esoteric and people who have collection problems may think, you know, I just have trouble getting through this stuff. That's so out in the, you know, weeds. Why would we ever think about that? But I think if you look at what's happening when you start collecting large data sets, which is essentially what's happening when we make digitized files of these things, the kinds of research, the kinds of relationships between these data sets that can happen in the future are things that we don't really envision right now.

I mean, once audio completely gets into that realm, and as Gene said, finish is not a term we necessarily talk about a lot, but as we get larger and larger data sets of audio, some of the things people will use them for will be non-audio reasons. They will be data sets that will be referenced and cross-related to other data sets.

And it's hard for us to imagine that, but that's, to me, what's really interesting about imaging is that there's a lot of things you lose once you, in terms of a record or a cylinder, put a stylus in that thing and only get the audio signal back.

There's a lot of conditional information, there's just the surface of these images have all sorts of information in them that we may not be envisioning as valuable right now, but are part of that record. And that's interesting because, you know, originally, when people talked about preservation, they wanted the original to live, and that's kind of the same thing again.

So it's funny that we're kind of going full circle that this technology is allowing us to what we originally would have wanted to do. You know, that doesn't say that this is the perfect solution. It's going to cover all, you know, it just wipes away everything else, that's not, you know, the case at this point, and it may never be the case, but it's something I think we need to keep in mind as we hear what people are doing, these projects that people are envisioning now.

How saving the image actually is a change. It really is a change in preservation. It's a fundamental change. So I'm really looking forward to this. You know, I've been working on this for, what, about 12 years now, Carl? So I'm extremely happy all of you are here and I would like to invite Carl up. I guess we're a little bit ahead of schedule, but that's not bad, is it? So, Carl.

DR. HABER: Good morning. So I'm Carl Haber. I work at the physics division, Lawrence Berkeley National Laboratory, which is part of the university of California, so this is an introductory, kind of, overview talk. I will try to address general issues of modern, non-invasive, and data-centered methods of recorded sound preservation and access.

I share with Peter, the thrill to attend this meeting with so many people who have worked on and thought so deeply and worked on these and related subjects. I would like to acknowledge our host, the Library of Congress and the agencies which have been so key in supporting this work internationally.

So I'm going to talk a little bit about the technical and historical context of sound recordings and imaging-based methods, and then I'll try to define a little bit, because I know there's a whole variety of people here with different experiences, what we mean by these methods. I'll try to address the question, you know, why do we bother doing this at all, and then what does the field, from my perspective, need to move forward?

So here's a zoology of early sound carriers. You see a lot of different pictures of disparate objects. They're all forms of sound carriers that encode sound in a whole variety of different ways, so you should be struck with the diversity and the lack of similarity among these things. And that's what we face when we look, particularly at this early period.

A lot of different materials, a lot of different formats, shapes, sizes, structures, modes of recording sound, so the challenge is, how do we deal with all of this? So there is, of course, as I said, this great diversity, but in the service of simplification, I want to focus in on two general formats that encompass much of what is represented in this diversity and it will help motivate our understanding of the approaches to digitizing these materials.

So on the top, I'm referring to the famous wax cylinder of Alexander Graham Bell, and Thomas Edison, and others in the early period of sound recording, so this is a cylindrical object with a groove that, essentially, makes a helical pattern around the surface. And that groove encodes sound in a vertical modulation of the surface.

So the bottom of that groove goes up and down, and if you put a needle on it, that needle is following the up and down motion. So there's, obviously, a cylinder, a couple of cylinders, and this picture is a blown-up region of a tiny, tiny portion of one of those cylinders, and so the groove is here between these two ridges, and time is in this direction, so you can imagine a stylus riding up and down.

This is a picture, it looks like a photograph, but you should actually think of it as a map of heights. So if it's dark, it's deep, and if it's light, it's high, and so as you go along the time direction, you're going deep and high, and deep and high, and so forth, and if you follow this, it's, basically, an up and down pattern.

And the extent of this motion is just about 10 microns, so that's not going to be a familiar unit to everybody. So think of a human hair, according to Wikipedia, the average human hair is 50 microns, okay? So this is a fraction of a human hair, but in that tiny amount of space, the audio, which has been recorded on the thing, is encoded.

If we look across, now you're going

around and around, you see that it's a kind of circular-shaped structure and the bottom of it is, essentially, the entire structure is moving up and down. So we refer to this as a vertical recording. And typically, cylinders are made that way, but there are vertically recorded disc formats as well.

Over here is the more familiar disc, the phonograph record, if you like, was the common term that maybe we used when some of us were kids. There, instead, you have a groove of fixed depth. For example, it might be around 75 microns and it's sort of like a 45-degree triangle, and the sound now is encoded in a side-to-side movement of the groove.

And so here's the needle and you can see in this picture, the groove undulating from side-to-side. And we refer to this as a lateral, or side-to-side, recording. And there are many variations on all of this, but most recordings can be classified as vertical or lateral.

Again, the size of the movement from side-to-side is as small as less than a micron and as big as many, many tens to hundred-something microns. Okay. So what about these -- what are these non-invasive and optical approaches to playing recordings back?

So you really start off, when you think about using light instead of a needle to get information off of a recorded sound carrier, back around 1960, people were already talking about using the reflected pattern of light bouncing off test records, records which just contained a set of fixed frequencies, as a way to do quality assurance and test these things.

In the mid-1960s, people started using the scanning electron microscopes to study the structure of grooves on records, again, as part of engineering. In the 1970s, interferometry, the interference of light as it is reflected from the surfaces, was used to study CD-4 discs. And then in 1977, people really started to work with the idea of using light to really play records back, not just to study them.

And in 1977, Heine patents a laser disc player. Okay. In the 1980s, a company called

Finial tried to take this system commercial and failed. It was then taken over by a company called ELP-J in Japan, and they turned it into a marketable product in the 1990s and you can buy these things now. It's the laser turntable.

In the late-1990s and early 2000s, a variety of, essentially, research labs and universities, and places of that sort, tried to use a variety of laser reflection methods, particularly to try and read cylinders. They're sort of small projects. Syracuse University and Lausanne in the Soviet Union and in Japan, so the names of the principle players are listed here, and there probably are some others that I neglected.

But fundamentally, all the methods that I just referred to replaced the stylus with light. So you're essentially -- you're still playing these things in this kind of continuous way, light is bouncing off a surface and being picked up instead of a stylus. Okay?

But the approaches that are really the focus of this meeting are what I would call metrological approaches, okay? Instead of replacing the stylus with a beam of light, the idea is to treat the entire surface as a very, very high resolution digital data set, a big digital data set, that you can analyze offline to extract the recorded sound in terms of a mathematical algorithm.

So we're not making a player anymore, we're making a way of migrating a physical object into the digital domain, and then applying large-scale data analysis tools to extract the information from it. And I think the originating idea in the sense of the literature, as far as I could find, is in a paper by Stanke and Paul, which was called 3-D measurement and modeling in cultural applications, which was published in 1995.

And they point out that a whole variety of materials, statues, ceramics, and all sorts of things could be digitally rendered with 3-D scanning methods and that would add value for scholarship, and museum work, and archiving, and so forth, to have these objects in the digital domain, and I know they were interested in sound recordings as well as statuary materials and so forth.

In 2001, Cavaglieri, Babst, and Johnsen in Switzerland developed a 2-D photographic method of transferring disc records into the digital domain, which is called the visual audio, and you're going to hear from Stefano here today about visual audio. And in 2003, a 2-dimensional and 3-dimensional direct surface metrology approach to digitizing these materials started to be worked on by our group in Berkeley and colleagues at the University of Southampton in the U.K.

In the U.S., we refer to this as IRENE. You've probably heard the term in the tours and so forth, so that's kind of where that originates. So what's the basic process? I want to transcend the specific limitations so I picked a picture that nobody will confuse as a sound recording, but the point here is, you see the David and this rig is some kind of a digital scanner that you can see this little patch of light on his nose. It's actually digitizing.

There is a digital David that exists

through this work and it's a tremendously interesting project, so think about a generic scanner and a generic object of cultural value, and then replace it with a phonograph record, disc, or a cylinder, and through some scanning methodology, and you will hear about different ones, we end up with either a 3-dimensional rendering of the surface or a 2-dimensional rendering of the surface that then is the basis of this digital data set.

Archive it and then you also pass it through an algorithm that calculates how would a needle move through this surface if it was actually playing? That creates an audio waveform which you can then put into the form of a file. So in the process of this, of course, you can do digital restoration, you can ask questions about the condition of the material, and so forth, through these analysis processes.

Okay. So the two basic ways of getting information once you're on the computer form these data sets are as follows. If you have a very high contrast 2-dimensional image of the surface, that can be suitable for grooves that move from side-to-side. You can obtain these either through direct imaging or through photographic films of the surface that you then scan.

You have to have an illumination system that highlights the surface features that are flat so you can see distinct edges. You need sufficient resolution to match the audio content, so that turns out to have about one pixel in your imaging system needs to refer to about a micron on the objects surface.

You naturally end up sampling the audio in the time domain up to 100 kilohertz, or even more, quite readily when you do this. There are high-speed cameras for collecting pictures like this that allow you to collect the data in just a small factor times the actual real playtime.

And then finally, you can extract the groove information and information about damage and debris by applying a very well-known mathematical image processing technique, which is called edge detection, and that's what happens when you apply an edge detector to this picture. It throws everything away except the interfaces between light and dark where all the information is stored.

Once you have this, it's basically a swarm of data points that reflects very directly, the stored audio information. It's almost trivial to go from this kind of a data set to actual sound, but the great thing about this is that it so nicely captures the very regular but undulating structure of the groove, but when there's a damaged area, or dust, or a big scratch, the edge detector kind of goes crazy and you end up with a swarm of relatively uncorrelated points.

And that's a great way of identifying damage and spurious features in the audio right there in the image. So it's a pretty powerful method and it's widely used among the different approaches.

Instead, if you go to three dimensions, you want to capture the entire 3-dimensional structure of the surface, that's important in the cylinder or any vertically cut material because photography is not a very good way of capturing the third dimension.

So the method of choice which has emerged among people working in this field is a type of confocal microscopes. It's a microscope which actually measures the third dimension very well and the other dimensions in a scanning kind of approach, and it's actually very clever the way it works.

It takes white light and it focuses it through a lens that's a mixture of a lens and a prism. So all the different colors get dispersed by this lens and they come into focus at different places. Okay. So you get this swarm of points of red, green, blue, and so forth, and then when you pass the surface through that region of multiple focus points, you get reflections back strongly from the color that's in focus and weakly from the colors that are out of focus, and so you can relate color to depth.

These are extremely powerful devices and they have a very significant industrial and manufacturing application, so it's a very growing field and we've really benefitted from these things. The point that they measure is very small. It's only about 3 microns in diameter. And they have vertical resolutions of 50 to 100 nanometers.

So 100 nanometer is 1/10 of a micron. A micron is 1/50 of a human hair, so it shows you, you know, what we're talking about, sizes which are small compared -- resolution which is small compared to the natural scale of how the surface of these materials change due to the audio.

And they're very verbose. You get a lot of repetitive data about the surface, so you can apply averaging and so forth. Okay. Now, when you take a physical carrier and you migrate it from the physical object to an electronic form, in the modern day, we're digitizing. Okay?

So it's important to understand what we mean when we say digitization, because it really has a kind of different meaning in these image-based approaches than in the more traditional approach, so let's consider the traditional approach.

So in the traditional approach, you put

a needle down, the needle follows the groove, the motion of the needle transforms -- well, the velocity of the needle is transformed by the mechanism in the cartridge into an electrical signal, that's what this waveform is, that's passed through a filter, and then into an analog to digital converter that's running at a certain speed, let's say, 44.1 kilohertz for standard CDs or 96 kilohertz for the higher speed archival specs, and that results in a series of discreet points that represent the filtered version of what was on the disc.

So there's a clock and there's an ADC. Now, when we digitize a sound carrier using these optical methods, we're essentially taking a picture of the surface. And so we're turning the physical object into a grid of points, so pixels, or voxels, right? Now, those pixels have different meanings in different directions.

So the pixels that are going along the time direction, the direction of increasing time, the way those pixels are sliced up tells you how fast you're sampling the audio, so that's like the clock, right? The pixels in the other directions, in some way, relate to the accuracy with which you're measuring the amplitude, but that's a much more complicated relationship, and it depends on many things.

It depends on the properties of the imager, it depends on the algorithm that you're using to get the data out, so while in the case of an ADC I can tell you, well, according to the specification book, this is a 16-bit ADC, or a 10-bit ADC, there's not a simple way to say that a particular optical process results in a certain particular bit depth.

There are many particulars, and I'm not going to address this point more until I get to my own presentation later about the IRENE system, and I will talk about it in that particular context, but I'll confide, obviously, to the other speakers to address that issue as they see it.

Okay. So why do this? And by the way, can you let me know when is this talk supposed to end? Keep going. Okay. So never. Okay. So why do this? Why are we doing this? And of course, Peter and Gene, I think, and Mark, eloquently addressed this, but to reiterate, we do this for preservation to protect delicate or damaged objects from further degradation, and to restore the unplayable.

We do this for access, to use automated scanning and analysis methods to massively digitize large collections. We do this to assess the, in a detailed way, the condition of these sound carriers. We do this to avoid the need to maintain legacy playback systems and diverse legacy playback systems.

And we do this in the hope of applying high resolution methods to extend the frequency response and noise reduction opportunities of getting data from these carriers. And I could add, you know, Peter's really nice point about getting access to what we maybe even consider to be intangible sometime; these hidden, latent sets of information that can be very meaningful in the end.

Okay. So where are the advantages? So remember that we're transforming the object now into a very large digital data set. This is a modern notion of the digital humanities. Okay. And once you get into the realm of the digital humanities, you can ask questions and do things that you never would have thought you could have done before.

So there are advantages and I will have a couple of slides that address these, but just as a list, I'm going to argue that this approach is very general. I try to make the point about the diversity of sound carriers. The methods are redundant, they're verbose, you get a lot of data that has repetitive information in it. They have extraordinarily good frequency response and resolution out to very high frequencies.

They allow you to deal in a rational way with delicate and broken materials. And they lack the physical dynamics of a mechanical playback system that responds in real time and that is an important thing as well. So generality. These methods have already been demonstrated on a large variety of materials and formats that required, actually, very little reconfiguration of the scanning system.

So shellac discs, lacquer discs, aluminum discs, Memovox discs, dictation belts, wax and plastic cylinders, copper galvanos, tinfoil cylinders and discs, experimental optical discs, paper tracings, wax discs with both vertical and lateral cut. These are all examples of things that have been scanned and restored already by optical methods.

I reproduced the picture from the first and second slide because that is actually examples of all things that have been already looked at to this day with optical methods. So sound redundancies. The sound is recorded, not just in one point on the surface, but in the entire profile of the groove.

The entire profile contains multiple copies of the same audio information, so the groove is moving, and the stylus, when you plop it in there, really only samples a small portion of the groove, but a more complete data set, the entire groove profile, gives us analysis options, which can sometimes add value.

And that's represented showing that these lines are the computer's choice of using all the points along the edge as an estimation of the sound and here's a case where there's some sort of a defect or a damage in this particular portion of the groove. That can be eliminated from the data set because it's so redundant.

So frequency response, I'm going to dwell on this for a moment before finishing up, because this is, I believe, a really critical issue for optical restoration and something that people really need to keep in mind. On a lateral disc, and when the groove is moving side-to-side, as the frequency goes up, what happens is, the stylus sort of gets pushed higher up into the groove.

On a vertically cut disc, something really different happens. As the frequency goes up, the stylus starts to sample the groove incompletely. And I'll show you in the picture why that is. This leads to distortion, to attenuation, and to non-linearity in the playback. Optical probes are so much smaller than physical styli that they don't suffer from this until you get to a very, very high frequency. Okay?

So these little undulations here are supposed to just, you know, be higher and higher frequencies that might be vertically recorded. So when the stylus is placed in there, you can just see very clearly from the picture, that at a certain frequency, the stylus is going to be too big to fit into the structure anymore vertically. Okay. That's called a tracing error.

Once that happens, the amplitude gets attenuated, the sound gets distorted, and it happens in a non-linear way because it samples one direction of the swing, but not the other. The optical stylus is so tiny that this problem only shows up at a much higher frequency. So where does that happen? So here's just a simple little calculation. Let's look at the numbers.

So a cylinder's playback stylus of 7.5

mils is a kind of standard value. Plops down into the groove, which, let's say, has a depth of 10 microns, and it reaches the bottom, and at some point there's something I'll call the critical chord, it's the place where, essentially, you just fit, okay?

So that occurs around 85 microns for this geometry, right? So the question is, where does the wavelength on the surface actually hit 85 microns and the cylinder playback stylus no longer fit anymore? Okay. Well, that depends on the speed with which the record was recorded, whether it's a 90 rpm or 160 rpm, so, you know, I did this whole spreadsheet.

But you can see that this stylus will start to not fit anymore at 3000 to 4000 cycles per second, okay? The optical limit occurs around 100,000 cycles per second, and it's just because one is bigger than the other. So really, if you play a cylinder back with a stylus and you digitize it at 96 kilohertz, you're effective bandwidth is equivalent to only sampling it at 8 kilohertz, which is twice that critical frequency. That's the

Nyquist Rule.

So people think, well, I'm running this thing at 96 kilohertz, or 48 kilohertz, you're not getting that information because the stylus just can't pick it up. It's too big. Who cares? Right? Who cares about high frequencies on an acoustic recording? Acoustic recordings only capture sound up to a few kilohertz anyway, but the problem is that damage and wear don't respect these limits.

Cracks, scratches, the effects of mold, have high-frequency information in them which doesn't care about the frequency limitations of the recording process. Okay? That's the reason, in any case, that you want to transfer a recording with very high sampling rate, because you want to get at the noise.

Only an optical scan gives you the resolution to measure and process these high-frequency features, okay? And here's an example. So this is a cylinder which was played back on a modern stylus player and was played back optically. This is a plot of the energy, or the power, essentially, at different frequencies.

So here's the music, it's the Rakoczy March, so if you're a Hungarian nationalist, you're going to like this, and once you get above a few kilohertz, the frequency content drops really fast. That's because the stylus is too big. All right. So let's listen to what it sounds like.

Okay. And now the same cylinder, the same portion, played back optically. Now, you might not like that because it sounds brighter, and tinnier, and is more noise, but that's exactly what you want. You want all the information, you want all the noise, because then you have what to work on when you try to process it and handle it, okay? So this is a more true rendering of the information than that which limits you when you have a stylus that has a natural frequency cutoff.

Here's where it comes in when you try to do noise reduction. This is a depth map, so dark is deep, of a little portion of the surface of a wax cylinder from the collection of the University of California's Hearst Museum. It's a recording of

Ishi, who was the last speaker and surviving members of the Yahi band of Northern California.

He was recorded in 1911 by Professor Waterman and Kroeber. A very famous collection. It's been on the National Sound Registry and it has damage due to mold that eats into the surface, so that's all this black stuff that is deep into the surface. It's dark. Okay? So when you play it back with a stylus, this is what it sounds like.

When you do an optical version and you capture all these fast noise transients and can process them, you're going to hear this. And don't pay attention to the rumble. I want you to pay attention to the crunching sound due to the mold damage. Okay. So broken and delicate materials. These are obvious for non-invasive methods.

A number of broken and damaged things have been shown already. You're going to see more later in the conference. There are issues related to how you collect data from these segmented objects and there are different approaches in visual audio, and IRENE, you're going to hear about those, but there are many common issues in the way the data has to be analyzed and how you link groove segments across gaps, and then the problem of, what happens when the number of broken segments approaches infinity? How do you deal with that? It's a very interesting problem.

I think finally, modeling. Okay. What you see is what you get. That's a little bit of a pun, I guess, but what I'm saying is, because you have these very, very high frequency response and very high resolution, you get optical transfers that are truly flat in the frequency sense. The optical measuring processes don't have an intrinsic frequency response which it then imposes on the measured audio.

A stylus, on the other hand, is a dynamic mechanical system and that creates a particular sound when you hit a defect. Okay. Physical modeling can be used to add these to optical measured data, but that's really a choice. So this is a stylus playback and you can see, it hits a scratch, and the stylus actually rings for some period of time.

A mechanical stylus respond dynamically to these imperfections and that creates the sound that you're familiar with in a stylus playback. But that's an effect, okay? It's an artifact. It's not actually part of the sound recording.

Is this a universal solution? No, not yet. At present, the tools are expensive. They're scientific instruments with a limited expert base. They take longer than traditional playback methods by a fair margin, in some cases, to collect their data. For commercially pressed shellac discs, which is most of what's out there in some numerical sense, and if they're in reasonable condition, traditional methods are faster, much faster, and are often superior.

But for vertically cut records, and delicate, damaged, or special needs materials, like lacquers, there are very, very significant advantages, and I think that's what brought everybody here today. So what does this field need? The field has made, I think, very good progress in about 15 years, considering the small number of people at work.

In the 1980s, over $20 million of venture capital, in then year dollars, was invested in domestic R&D for the laser turntable that actually never really worked, and then the Japanese started working on it. So I can only imagine what they invested in it before it became a commercial product.

The amount of money which has been invested internationally in optical methods, of the meteorological sense, is much, much less than this. We need a larger community of developers and users to bring this into greater use. We need to develop accepted standards and specifications that everyone agrees upon and adhere to.

We need more targeted engineering and software development. We need access to new instrumentation as they appear. We need to offer more training opportunities, educational opportunities, and internships to younger people to learn these methods, and we need more focused large transfer projects as ways to gain experience and converge on best practices across the community.

Okay. So I'm going to conclude just with a little anecdotal reference. When you're confronted with an opportunity, a new idea, or some direction to go, obviously, people are often hesitant and maybe don't want to take the chance to try something new. I think the Library showed a lot of leadership and the funding agencies in Europe who supported this showed a lot of leadership to embrace these things and say, you know, why don't you try to do these things?

But you should always keep in mind that sometimes your assumptions, if you're scared of them, they prevent you from going forward. So I want you to question your assumptions. So the invention of sound recording in the 19th century utilizes technology and methods which actually could have been applied hundreds of years before that.

So I always wondered, could Leonardo da Vinci, who was, like, the smartest guy of the 1400 to 1500s, you know, have invented sound recording? Why didn't he do that? He had everything he needed to do it. But it turns out that da Vinci believed that poetry and music were inferior to sculpture and painting because hearing is, and I quote, less noble than sight, in that, as it is born, it dies, and its death is as swift as its birth.

Okay. So he figured, you know, God didn't care about sound or he would have made it more permanent. He didn't realize that he was put on Earth to solve that problem and find a way to record sound, but he just didn't do it. Okay? So, you know, it's always good to think about Leonardo when you're trying to make your next funding decision.

Anyway, so I think that's all I have. Yes. So thank you very much. So I think you're doing questions at the end, right, Peter? Or now?

PARTICIPANT: So, Carl, as you're talking about the maturation of this, I guess, traditionally, if I went to the Library of Congress to get a recording, I'd get a recording, but now, is there a method to distribute the data sets? Is that the future you're seeing?

DR. HABER: Well, so you're asking a question about how the results of these optical studies are disseminated, okay? So for small pilot projects, and I'll speak about some that we did at Berkeley, and I'm sure our colleagues will speak about other ones that they have done, we have setup, you know, Web sites about the Volta Laboratory collection from the Smithsonian, and so forth, but in terms of the Library's process for disseminating the results of this, I would like to defer to Peter or colleagues from the Library to answer that question.

The question is, if you do these large-scale transfers, how are the results made available to people?

PARTICIPANT: Can I get the data set or do I just still get the recording?

DR. HABER: You mean, you want the image data.

MR. ALYEA: Well, I mean, the mission of

the Library is to make this stuff available. There's nothing -- you know, if this was previous to the ability to image these things, and someone wanted to physically look at an object, you know, they'd have to come here and see it. That's similar to being able to distribute a data set. I guess there are copyright issues, right, Gene?

I mean, if something is copywritten, I suppose the image data set is as copywritten as an audio file, so you certainly have that. You know, you can't just pass these things around, but do you have something --

MR. DEANNA: Well, you're talking about doing something that's completely new, so, you know, taking copyright aside, if you were asking for, you know, an IRENE transfer, say, of a sound recording, in my opinion, you would get the full data set and an audio file. That would just be what you got, but that's a new idea, although, you know, now, we'd provide you with maybe a scan of the label. People sometimes ask for the scan of a label of a disc and the audio file as well, but nonetheless, that's a new idea, but yes, I would consider that part of the transfer.

And I think NEDCC, who is actually doing this work, does the same. I'm certain that they are distributing the data set along with the --

DR. HABER: Yes, so I don't know, is Bill or Mason here?

MR. VANDER LUGT: In general, for institutional clients, we include the image data, and some of them aren't interested, but we've had some people that were more interested in the images than the sound. Yes, I think they're really important and used a lot.

MR. FABRIS: My name's Jerry Fabris. I'm the curator at the Thomas Edison National Historical Park and I can speak to this as a user, as a customer, both with experience at Library of Congress and at NEDCC, and I think with this process it's important to understand that the archival data that's captured is the image data, and we, in my experience, I was given the image data, but no capability of using it, not the software, in other words, so I think that's something that needs to change.

DR. HABER: I think that's a tremendously good point, and so what you're saying is, to translate to a more general way, so the optical systems scan, they create image files, then the operators use software to transform the image into audio. And so simple-minded thing is, okay, I'm going to give the user the audio. If you give the user the image, there is software that can be used to transfer that.

And you can make different choices when you do that, about noise reduction, or bunch of choices that you can make, and one of the goals really should be as part of this notion of larger community of developers and users, and more targeted engineering, and software, it would be great to make an application that is easy to use, that is very user friendly, that has the basic presets or choices, and it gives a person with just a little bit of training the possibility to process the image data, and it's a great goal.

It's just something that hasn't been accommodated yet, but it's certainly something that you could define quite readily and think about doing. And you're right, that is something that you should have.

DR. NYE: Jim Nye from the University of Chicago. Thanks, Carl. You use the trope of zoology, but it occurred to me now as you're talking, as this discussion is going forward about images and distribution of images, that you've added a new element to the zoo in an important way, and extended the zoology, really, carrying it forward.

But speaking to the importance of distributing the images themselves, if I think of an analog, the Dead Sea Scrolls come to mind. And for a long time, many of you probably will know that it was only the research on the Dead Sea Scrolls that was distributed and the images themselves were not, until fairly recently.

But by virtue of distributing the microfilm or other sources that have the full context of the Dead Sea Scrolls you've increased the number of users. And here, I'd like to point back to your set of needs. You've talked about a larger user base, but presumably, you want a larger resource base as well.

By distributing the images, effectively, you increase the possibility that a larger group of scientists can address these questions of interpretation of the signals in new and creative ways that might not be possible now or even imagined now.

DR. HABER: I mean, it might be a good idea to sort of do a pilot project where we put up a variety of image files from a variety of sound carriers, and put it out there to the world of universities and interested people, like a data challenge, to see what they could do.

I actually went to a data challenge event in San Francisco. This where these big data companies offer a prize to come up with a better was to analyze things. So for example, I'm a high-energy physicist, and they put out the data set that led to the discovery of the Higgs particle, and they just asked the world to try to analyze it and see if they could get a better signal-to-noise than the physicists got, they did, but these kind of data challenges, you know, they gain some traction.

And there's a lot of smart people out there who think about large data sets and so it might be a way to seed something like this. Yes, there's going to be a general discussion after.

MR. ALYEA: Okay. So Ottar Johnsen's going to talk next.

DR. JOHNSEN: Okay. Thank you for the opportunity to present the VisualAudio project. I would like to say that many people, same as with Carl, have contributed to this project. Some of the names are here. I have forgotten many and also, there would be maybe 50 names to add. I am going to jump over some of my slides because Carl had already given you the most important information, that's why I go to next slide.

This slide is mainly interesting because that's the kind of slide we have been working on, and Stefano will get into more detail about such records. For Carl now has talked about the fundamental, and now I'll go into the fundamental technique. It's an idea from Stefano, I will talk a little bit about it later, is the entire surface of a disc could be photographed.

Here is the same observation if we look with a microscope of a groove, probably, you know what kind of record this is. This is a stereo record because you see the width and the lateral movement change with time. And here is a basic block diagram of the VisualAudio concept where we have one extra step. From the record, first, we make a film, then we digitize the film, then last step, we go from the image to extract a sound.

Now, why should we have an intermediary photographic step? It's controversial. Some think it's a good solution, some think less. I believe one idea is that it's an urgent way to save a disc before it decays even more. Film are small, cheap, quite stable. We know it, thanks to microfilms. Microfilms last a few hundred years, it is supposed, and we are using film relatively similar to microfilm.

Here is what we get if we illuminate a record from above -- both with a 78 rpm where the groove is a little bit round and with a 33 rpm record where the groove is a much more triangular groove. And this is what we get at the bottom of the groove on the film, but I will get in more detail on it later.

So what do we need for the film? Sure we want the highest possible resolution, the smallest grain, it should be black and white, we don't care about color, and also black and white film has much higher resolution than color photographic film. And also, the film speed should be high so that we can reduce the time, and also, we can reduce the amount light.

We melted a few records, I believe, at the beginning on our first trial because we had 500 watt light, but we did it on record of no value. Here, you see the picture of the system. And here, in the next, we see the system to take picture.

Basically, the system, where we make a 1:1 picture, we have a 40 cm focal length lens, if we do a 1:1, we need 80 centimeter (two times the focal length) on each side of the lens, so it gives a total length of 1 meter 60, and the total system is more than 2 meter because we need some space for mechanical hardware outside.

So how do we make illumination? We make the illumination so that we can separate the interesting parts. We use directional light, we use monochromatic, or nearly monochromatic light, so that, first of all, we select the shortest possible wavelength, and that is blue, or deep blue, and secondly, we have less problem with chromatic dispersion when using a single color.

And now in the photographic part of optic, resolution is limited by two main factors. We have the diffraction that gives it an Airy disc, it means one point become a spot of a certain size, and we have the de-focusing error, because we are never completely focused, just because our record has a certain depth.

And both of them give a certain spot enlargement and we can calculate all of them, and they depend on the numerical aperture. And both of them depend -- and one of the problems is that one improves when you increase the numerical aperture, the other one go down. So we must get a compromise. So the good thing with a compromise is that when we have the sum, the bottom is very flat.

So we can use practical consideration, for example, the size of the system, to select something that is within 1 or 2 percent of the optimum. Here are the values with the end number that we can practically have on the optic we had. And here again, you see the resolution in micron. So now we have done one process, we have taken the picture, and what's interesting with the picture is that, with this picture, we have a big depth of field. We have a depth of field of 1 millimeter.

And that is good since many records are not flat, are in bad shape, can be bent, but a depth of field of 1 millimeter. We can then get bigger depth of field, we have less resolution.

Now, next step, we have a film, we use negative film, but we could use, also, positive film, and then we have to use a scanner. And our scanner is quite similar to Carl's 2-D system. Simply, we are scanning a film, not a record. In place of having reflective light, we have light going through the film by transparency.

And here, we see a little bit better, you recognize on the bottom of the system, a light source. Unhappily, we had to use red light because blue light source are not available. It is LED, actually, already. And we have a rotating scanner. At each rotation we have 2048 samples. So we are scanning in each rotation, a certain number of grooves.

Here, there are only five, but usually it's 10 to 12 rotations we get when we scan. Our sampling frequency is variable. We have used sampling frequency of up to 260 kilosample per ring, corresponding for 33 rpm record, for sampling frequency up to 110, and for 78 rpm record to a sampling frequency of up to 260 kilohertz, but we used less, but again, as Carl said, out-of-band noise can give useful information to reduce in-band noise.

So here is what we get when we scan the picture. That is a more typical example where you recognize the sound and you see that for each groove, we have two white line and in the middle we have black line. Remember, it is a negative. So with all we talk about, the light get inverted.

And now, if we look at the optic, we go from the photographic optic to microscope optic, we have the same physical constrains from the optic, but different number. We have the diffraction, given the Airy disc, and we have the de-focusing. And there, we consider that we need a depth of field of 30 to 50 micron, and it give us blur, total blur, of about 13 microns.

And what is interesting is, as it is shown on this curve, increasing magnification does not increase the resolution because when you increase magnification, at a certain point, you have a reduction in the resolution. We have different figure whether we need a depth of field of 30 or 50 micron.

Where do the number 30 and 50 micron come from? They come from our knowledge about the negative film, the inaccuracy in negative film, plus, the film is put on a glass plate, and the glass plate rotate. We have measured how the glass plate move up and down on the rotation, and that is how we came to the fact that we need a depth of field of between 30 and 50 microns.

And here, I go back to what is happening with 78 rotations per minute record, we have a form of the groove, and it gives, where we have a lot of light, reflected light, it's black because it's negative, and where we have no reflected light, it is white, and that's why we have two white lines for each groove.

For 33 rpm record, usually we have nothing at the bottom, but sometimes, because the cut is not perfect, a very thin line, black line, appears at the bottom.

For processing the image, we do,

basically, as Carl said, edge detection. And we have two or four edges. Secondly, important thing, the edge detection can be detected with better than a pixel accuracy, what I call subpixel, accuracy. The accuracy of the position of an edge depends on signal-to-noise problem, not to the size of a pixel on the scanner.

And as I said, noise and distortion are the problem, the main problem, it is not the blur. It is mostly in the noise and distortion. But if we have less blur, we have more information on the noise, and we might be able to get rid of the noise in a better way.

And the noise distortion, there are many of them, dust defect, non-constant illumination, our main noise that is different from what Carl does, is the film grain. We have random shaped crystal in the photography. We have noise naturally in the linear CCD camera. We might have, but I believe it is very low, our sampling frequency of sampling time is not perfect.

There might be vibration, mechanical

vibration, on the scanner. There might be bad centering of the record. Same issue. And bad centering create only second or third order problem.

We can't have constant illumination due to the optics of the microscope. That can be corrected with a curve, but any correction naturally introduce some slight noise. And it's easy, we measure the film without anything, and that give us a correction curve.

And now we look at the processing. Here is a typical example. We have two white line and we want to extract the edge. And we are looking, we cut through it, looking at one pixel and you have intensity. We detect the constant level on the top and the bottom, and after the threshold is put about in the middle. It is more complicated than that.

And then, we are going to do it with time, yes, vertically, it's the time, or it is space here, but it corresponds to time. And now, what are we going to do? We are going to look at the image to find mistakes. As you can see below, if we take just the average between the four edges, or between the two edges, we see clearly, some distortion.

And we have some information to get rid of distortion. We know that all the four lines should be parallel. So we are being stupid and adding the four line to get the best position. We add the four line when everything is perfect, or nearly, but when one line is clearly different from the other, we suppose that that line is incorrect. And then we will make an interpolation on this line based on the three other lines. And when I say line, I mean, naturally, the edge.

And typically, also, we use co-width variation. And when we have cut, we do interpolation, linear, or first or second order interpolation. And here, in another example, where you see where we have the four edges, and where we detect where there are problems. And in red, we have the edges, the four edges, after correction.

Here is, on top, we have what we -- the basic signal, the edge signal, and on this edge signal, I like to show the top one, we have two special effect. We have a ramp. The ramp is due to the spiral in the groove. Then we have a sinus, nearly a sinus, wave. The sinus wave is due to bad centering of the record.

And then, at the beginning, you see some jump. That's when we jumped to the next groove. We made a mistake. We jumped from one groove to next and then back again later. It didn't happen much, but on purpose, I put one bad case where we put it in evident.

So after taking away the ramp and the sine wave, and the sine wave has a frequency equal to rotation speed, for 33 rpm records, it corresponds to about a 0.5 hertz signal, for 78, about 1.2 hertz, about, so it is very easy to get rid of it just by high-pass filtering.

And then afterwards, we get the sound with glitches placed where it jumped or where there were some noise. Afterwards, I don't want to go into all the specialty record, but most record are equalized. It means that what you see, the lateral movement,is not the sound. There are some derivative and integration to be made before getting the equalized sound.

It means that when reading, we must make inverse equalization operation as the one done when writing on the record. In most mechanical playing device, the same is done. The problem, we don't always know how the equalization was made. We know it for after the 1930s. We had standard equalization. Before, each record was nearly equalized in a different way.

And additional post-processing might be applied, but many sound specialists say, don't process, don't do additional post-processing on the original. Only do it on the copy. Keep the original, even in bad shape, because later, we might have better technique.

Now, how can we measure the quality? That's a big issue. Perceptually is one possibility. Another possibility is to measure the noise in silent section of the groove. A third solution is to have a test record and to measure what we obtain. As an example, this is just for showing, on the bottom side, we have the spectrum of a sine wave.

And what do we have if we have a perfect sound on a record? We'll have a peak at the frequency of the sound, and elsewhere, we will have noise. That's what we see at the bottom. And at some high frequency, you have some small peaks that can be due to many phenomena. For example, 30 or 60 hertz interference during recording that could come out a different frequency later.

And on top, we have an example, that is what happens when we have non-linearity and we have an error that creates harmonics and inter-modulation. So such techniques can be used to try to measure the signal-to-noise ratio.

And what resolution do we need to get, for example, a 40dB signal-to-noise ratio? With constant amplitude record, we need a resolution at (or accuracy of) 1.75 micron. For constant velocity, 0.25 micron, and for equalized record, 1.28 micron. And this is amazing number showing that in mechanics, 150 years ago, they did as well as electronic managed to do with silicium 20 years ago.

So electronic is more than 100 year behind mechanics. I like to tell is, especially to my electronics colleague. I am electronics also, no mechanics. I show how perfect one was able to record things in a mechanical way. So does it make sense to try to get this? No, a blur of 25 micron and a resolution of 1 micron, I mean, we are far away, but a blur is low-pass filtering.

So resolution is related to the noise. That's why it is possible, but naturally, if we can reduce the blur, it will be a good solution. It helps always.

The advantage, I believe I have talked about already, compared to mechanical and I want to go now into the detail of the lacquer record. The record cut in many pieces. When you look at the image file when you scan a record, the round record become a file like this and the sound is in a groove vertically. And we have big cut. And we can consider that we have pieces of the puzzle, except they are in position, but with a slight cut between them.

So what can we do? First thing, we must label each part of our puzzle as it is shown here. And now, we must try to fit, perfectly, the different pieces of the puzzle using information. And there are several parts, first, we have to detect the border, common border, of a chunk, then we must look at the best match, and then we must add together, or put together, different part of sound on each part.

So here, as an example, we have a chunk, and you see the vertical line correspond to the groove. And afterward, we have now, the common part between the groove and we must do the best match of the two blue part, or the two red part, or the two green part, and this is done as it is shown here.

First, we look at the parts of the groove near the border on each side and we look at what shift give the best match. We have found that there can be shifts equivalent to five groove, maybe even six or seven, so we have to test what displacement give the best match, it mean, give the best similarity in statistical property on the groove above and below the cut, and there are many different algorithms.

And there are some special situations. For example, you see here, sometimes one groove might disappear, that's what we see with the green, because it was exactly where there was a crack. Okay. And here, you see how we can see relatively often.

Now, what happened? In some case, the computer is much more stupid than human. So what did we have to build? We have to build something so that when the computer is not sure, an operator can decide. On the left, the operator for each place where the computer can't make a decision, the operator can change what groove fit each other.

For example, he can say that the groove should be fitted this way. There are still some problems. Sometimes the circular crack can be a problem. Secondly, sometimes, on record, there are several spiral. A record stops, and for example, because it was, for example, for a court case. At the end, they will lift the stylus and a new one will start.

And there might be some other problems. That's what happened with our first record, the reading was the wrong way. And with the first, six months after we started the project, we got very strange sound until we found the right idea, let's try to turn around the sound, so one should never forget that.

And then afterwards, simply, we concatenate all the sounds together to get the sound. And we have processed the number of such records. And typically, I just want to say, we couldn't recover all the records, at least at that time, typically, 35 faces out of 46 could be extracted. Some were in too bad a shape.

And now just a few more words. As a professor, I am very happy that a project doesn't finish by collecting dust because that's what happened with 99 percent of the projects of professor at universities, but the one who initiated it, Stefano Cavaglieri, from the Swiss National

Archive in Lugano, is also the one taking it over. And I'm very happy, especially that I am now one year before retirement age, so it's good to know that things are continuing.

It was very interesting because it's very valuable. That's why I didn't believe in this project in the beginning, but I knew it was a very interesting project for the students, because it includes optics, electronics, mechanics, everything, and what I found later, so very interesting, is working with people who are not engineers or non-physicists, but people from the humanities, and people from, yes, the cultural world.

And it was collaboration between many institutions and with funding from different entities. Thank you very much.

MR. STORM: Hi, Stefano. How are you doing? This is Bill Storm. I have a question regarding the way you've -- in the 1980s, there was a major project done in this country, the result was called a Rigler-Deutsche Index, and that's a microfilm copy of many thousands of sound recordings that allowed the Library of Congress, New York Public, Yale, Syracuse University, and I'm going to forget somebody, we did photography of all of these sound recordings with very, very even lighting, and it was using high-contrast microfilm.

Is it possible that your system might literally be able to play those images from the Rigler-Deutsche Index? They're very high resolution.

DR. JOHNSEN: Yes, what are the size of the films?

MR. STORM: 35 mm.

DR. JOHNSEN: 35 mm, no. That's what the first test, we tried 35 mm, and just by simple calculation, we found it was completely impossible and we arrived to the fact that 1:1 was about what would make it possible.

MR. STORM: Okay. That's too bad. Thank you.

MR. ALYEA: Another question for Ottar?

PARTICIPANT: I was surprised to hear

that your depth of field for the photo was only 1 mm, because it seems to me that warped records tend to be, or can be, warped by much more than 1 mm, so if you could explain that.

DR. JOHNSEN: Yes. Our depth of field is 1 mm, but simply, if we can go to much higher value, simply, than the resolution is lower.

DR. HABER: But warped records are a problem.

DR. JOHNSEN: Yes, but I would say, I made calculation showing we have, I would say, the order of 3 mm of depth of field.

DR. HABER: But warped records, which are acetates, probably, you don't, because of the glass or aluminum backing, get that huge wrappage where you might get with a vinyl because they don't have that same heat problem. And so to the extent that you've been focusing on lacquer discs, I think this large wrappage probably doesn't come in.

DR. JOHNSEN: No, I don't believe so, but Stefano maybe can answer later also about it.

DR. HABER: Yes, the delamination, so

that curvature is -- yes, but that's not a couple of millimeters, typically.

MR. VANDER LUGT: I think the warp can actually get up to 7 mm, if you include the warp of the disc base and flaking lacquer, but if I understand correctly, if you decrease the magnification, you can continue to increase the depth of field. So if you were willing to have a lower resolution image file and lower resolution audio, you could get a high depth of field that would allow you to image a delaminating lacquer disc.

DR. JOHNSEN: Yes. Or we can reduce the aperture. Yes, you don't change the magnification, you change the aperture, and that's the good way to get the bigger depth of field. Yes, what you pay is mere blur.

MR. ALYEA: Okay. So are there -- we're going to move on.

MR. CAVAGLIERI: Good morning, everyone. Well, that of Dr. Johnsen was the, let's say, scientific approach to this project, or to this system called VisualAudio; mine will be a very practical approach based on a day-to-day experience, as we use the system at our facilities at the Swiss National Sound Archives.

It has been a long journey because we started, effectively, with this project more than 15 years ago. It was back in 1999 and I'm happy that by now we at least got some tangible results. But of course there still are challenges, there are ups and downs and there is, of course, room for further improvements.

The basic idea, as Dr. Johnsen said, came from the microscopic observation of the surface of the disc. You know, if you look at the shape of the groove, it visually represents the acoustic vibrations which are very close to the corresponding electrical signal of the recorded sound.

And there is plenty of information available everywhere, but it requires a very, very high resolution photograph or scanning. This is just a schematic overview of the VisualAudio concept. I know you're getting familiar with that, but I want to insist on a couple of important points.

You know, what we do is: we actually photograph the disc, then we scan the film, and finally we process the image. But we've been talking about photographs, what does it look like? It looks like this. You know, this is an actual photograph shot with the VisualAudio system. It has a 1:1 ratio, which means, if your record is 25 cm in diameter, it's going to be 25 cm on the negative film.

So why do we have this analog photographic step? Dr. Johnsen already told you about a couple of important aspects. To me, having put this system into a productive environment, those points are very, very important. The first one is, well, it's a quick and reliable way to freeze the degradation of the surface of the disc. You know, you just shoot a picture.

Even if within the next few years, your disc is degrading, I mean, it's getting destroyed, you'll have that frozen image. We also have a certain depth of field which is granted by photographic principles, we've been mentioning that earlier on, and we can also think of the picture as some kind of an analog copy of the record, as it has a very, very, very high resolution.

If you look at the picture through a magnifying glass, for example, you'll see that it is just perfect. It's like looking at the grooves under a microscope. And film is small, cheap, and stable - while it's still on the market - I'm trying to anticipate some questions.

There are some additional benefits of the analog photography. Of course, the picture is shot without interfering with the surface of the disc. There is no need to manipulate the disc except for placing it inside the photo camera. We can take pictures of discs in virtually all conditions, de-laminated, broken, deformed, and they can, later on, be read and hopefully the sound can be restored.

Each and every disc format, it doesn't matter the size, the speed, the cutting, up to a certain limitation, is read using the same equipment. In some cases, well, under some circumstances, we even tried to get a picture of a vertical cut disc, a typical Pathé disc. I can tell you that we were able to look at the groove, even if it's just a 2-D system, a two-dimensional system, because the width of the groove changes in function of the depth.

So we were able to extract this information from a 2-D image. Image processing is well established, it's easy to make corrections to the physical variances of the disc. And again, fairly stable, small, cheap, for storing sound information, even if it sounds strange, you know, storing sound on film.

But this means that it might be implemented, or this is at least what we do, as well as a long-term storage format. So what we are doing right now, for example, is: we are taking pictures of some thousands of acetate records, because it doesn't take that much time and it gives us the possibility, when storing these things in our archives, to process them later on, so we'll still have time later to process these images.

You already saw the picture of the first generation photo camera. It was big. It was huge, not just big, and heavy. And this is where it is right now. I guess you can't figure out where it is. It is a sound recording studio in South Africa and this is the way we got it there.

One funny thing is that our colleagues at the engineering school in Fribourg had to literally cut the frame of the door of the room where the camera was built because it couldn't get out, you know? You can see it here. And well, you can see we had some troubles taking the camera into the studio in South Africa, but it survived, and now it's there. It's working. You know, nothing is broken. It's a simple thing, a simple setup.

Here, we look at the, as for the main parts, third generation photo camera. I'm not showing the second generation of the photo camera, which is the one that we are using currently at our sound archives. This is the third generation photo camera which is being built right now. This camera will be delivered in a couple of weeks in a photographic lab we've been setting up in another part of Switzerland.

You know, being the third one has the advantages of having fixed all the problems of the first and the second one. This one can be transported, it can be just, you know, disassembled and reassembled quite easily, and it's very, very, very easy to operate.

Then we have a turntable scanner, which is our kind of turntable where we put the film for reading it. It's still the first generation. It's the only one existing and this is probably the weakest part of the whole chain right now.

So as I said, practical VisualAudio, let’s have a look at the workflow. The workflow is a three-step process. Well, there's the first step, which is the photography. You know, photography takes, more or less, ten minutes for one side of a disc. Why does it take ten minutes? Because you have to position the disc, you have to shoot the picture, and you have to develop the film. It's an analog film, you have to develop it, and an automatic developing machine takes, more or less, five minutes to develop such a film.

Of course, if you have two operators in the same room, you can cut the time by two because, you know, one shoots and the other develops, and the other shoots, and the first develops, so you can really make it much quicker.

Then we have the film scanning process. Film scanning takes, let's say, 30 minutes per face, or per side. We have the film positioning, we have a ring-by-ring scanning procedure, and a merge procedure in order to acquire the whole image of the whole surface.

And finally, we have the image processing and image-to-sound transduction, which can take up to two and a half hours, where, you know, the groove is aligned and reconstructed, the image is transduced to sound, de-emphasis is applied, if necessary, and stuff like that. So the good thing is that the three processes are independent.

I mean, of course, you can scan a film if you have the film, but once you have the photograph in your hands, you can decide, whenever you want, to go further on. And that can be carried out by different people in different locations on a different schedule. The timing as for steps two and three is just a guidance. And of course, it relies very much on the shape of the discs.

When I say half an hour for scanning the film and two and a half hours for extracting the sound, I mean, if you're processing records like this, you know, completely delaminated. Things that you cannot read on a conventional turntable. From my perspective, the main, let's say, target for such a system is not to read records in perfect conditions simply because the quality is not good enough; I mean the sound quality.

Now, let's get more practical, I have some images and some sounds for you. First of all, we've been talking about rings. Okay, how does a ring look like? And this is in raw format. That's the reason why I have to set some parameters-- okay. So this is the ring. It's a number of grooves for one rotation.

So we scan a ring, then we scan the next ring, we merge the two rings together, and so on, up to the end. And this is how it looks like just after scanning, you know, working with raw files is not that easy, I don't know if you can see the green and blue lines on the edges.

So there is a green line here, a blue line here, a green line here, and another blue line here. These are the four edges of the two walls. I mean, the upper edge, the lower edge, the lower edge, the upper edge of the two walls of the groove. And these lines are the result of the first processing, as Dr. Johnsen said.

So this is the, let's say, how we estimate where the groove is and how we estimate how the groove moves laterally. Then, you know, linking the grooves is not that magic, you saw these pictures before, you saw that we can apply an offset to the image just to tell the system: when you come down here, you have to continue on that one instead of just continuing straight on because, you know, the two parts, the broken parts, of the record do not match.

And here, you can see a little bit better with a wider gap. But getting some sound is better, I hope, so I prepared a couple of different things. We didn't test the level, so I just play it now. Okay. That was read with a turntable by applying the emphasis curve that is necessary for properly replaying this record.

If we do a 1:1 comparison with what came out of the system when we got it from the school, it's this. From here, I don't know, I cannot make any judgements, because I don't know if you have enough low frequencies and things like that, but the two files, or the two sounds, are, in a way, similar. I'm not saying they are equal, they are, in a way, similar.

Let’s take another look at a different version, which is this one. This is a linear extraction with a turntable. And now the linear extraction with the VisualAudio system. We have some audio engineers here in the room and I'm quite sure that they noticed how linear doesn't mean linear when we switch from a turntable to an optical system.

This has different reasons, but it also sets some new challenges. And if we want to go into the challenges and start thinking scientifically, we may come up with things like, why don't we take a constant amplitude output, which is this? Or a constant velocity output. Or, what happens if we just apply a RIAA curve, a de-emphasis curve?

You know, I have no answers. Well, I might have some answers for you, but I'm not mentioning my answers. My main concern here, and I saw something similar, a similar situation yesterday, while visiting the Culpeper facility when we had a brief demonstration of the IRENE system, is that, when looking at those things from a scientific point of view, we get a lot of different options.

A lot of different options which do not mean too much to an audio engineer, or not always. An audio engineer is used to, you know, take the record, put it on a turntable, select the speed, select the correct stylus, select the correct equalization, and just do the copy, and that's it.

And the result, the end result, is absolutely correct, I mean, from an audio engineering perspective; it is what you want. With this kind of approach you have a lot of different options, you have the theoretical possibility to reproduce equalization and stuff like that as it happens in the traditional electromechanical world, but the result you get is not exactly the same. So why?

Let me take you to another example. This one, different piece of music, is a very early extraction. That was done using the VisualAudio system at the engineering school in Fribourg. This was just applying a low-pass filter at 21 kilohertz here. The same thing applying a 48 kilohertz filter is this. It has a little more high end, but what makes me uncomfortable (I should say, what makes me sick) is that the next version is NR. Do you know what NR means? Noise reduction.

So this was supposed to be the noise

reduced version of the previous extraction, which is not true. Why did that happen is because if I go here, you can see I have two folders. One folder stores the audio as I got it originally from the scientists, who not always have a deep understanding of audio, or maybe the knowledge, yes, but not the feeling, the right feeling, while the other folder, which is the one where the files, the sample files, reside now, is a folder with all files loudness normalized, I mean, not level normalized, but loudness normalized, just by applying the EBU R128 compensation, which is the equivalent of the ITU 1770 here in the States, which is what makes the perception more or less the same of a certain sound.

So if you take the original wav files as I got it from the school, they were just, you know, the noise -- yes, that one with the noise reduction had just such a low level that you couldn't notice the noise, but you couldn't notice the music either. So, we went a little bit further trying to measure things.

So we took a test record, a standard AES

78 rpm test record. You can still buy it. It's still on the market. We put it on a turntable, we extracted it on a 96 kilohertz, 32-bit float file, and got this result. And then we did the same thing with VisualAudio and we tried to match the loudness of the 1 kilohertz signal.

So, I'm not saying that the system is bad, I'm just saying that there is still some work to do, especially if I look at the noise floor, for example. So, this is the typical noise floor you can measure out of a turntable. You can just forget this little green line here. There's practically nothing. This is the same thing - yes, okay, thank you - this is the same thing when I measure the output of the VisualAudio system. This is the typical noise floor I get out of the system without applying any corrections, just linear. And of course, this has also an impact on harmonics. Somebody said harmonics are good, somebody said harmonics are bad, but we are used to harmonics.

So when you play an instrument, you never get a pure sinus out of it, you get the main sound, or the main tone, and you get a lot of harmonics out of it. So if we consider harmonics, I'll go back to that one, and switch back to the 1 kilohertz, sorry for that, but this time you see very clearly, we have a 1 kilohertz signal, a 2 kilohertz harmonic, a 3, a 4, which are visible. This is what comes out of a turntable.

And I'm sorry to say that again, this is what we are expected to get out of a sound system. If we look at the same thing here, we can clearly see the 1 kilohertz signal, more or less, we can also identify the 2 kilohertz harmonic, still there, but the others are gone.

So you know, this is my last slide, it says, well, may we get rid of all that noise? I'm talking now as an audio engineer. You know, sound engineering is my very first background. I've been working in that field for more than 30 years doing recordings and all kind of things.

You know, what I see is that… and maybe I'm wrong, but I would like that somebody in the scientific world would investigate these problems.

Things that bother me a little bit are, for example, the current image sampling frequency. We say it's 64 or 128 kilosamples per ring, while if we think audio, we talk about 44.1, 48, or 96.

How do these things correlate together? Do we have an impact in the audio quality because the two figures do not match? I don't know, but I would like to hear a very solid answer, not a theoretical answer, a very solid answer. On the image, the radial position of the groove can be estimated by detecting the edges of the groove, two edges, four edges, that's fine.

On a turntable, it is determined mechanically by the pivot and the inertia of the arm, as well as the stylus tracking the groove. It's a different approach, it produces different results. Can we measure that? Can we compensate that? Can we replicate that?

These are all questions I have for my friends in the scientific world. And maybe the last one is, you know, image processing extracts the position, in fact the amplitude of the groove, while a stylus gets the velocity. That was very well described by Carl Haber.

Is the derivative, which is the signal, the same for both methods? If so, can we apply the same de-emphasis, such as the RIAA or other less-known curves? I would like an answer, but not a mathematical answer, a practical answer, something that I can measure, or even better, something that I can hear.

Oh, there's another little question, is maybe IRENE suitable for processing our analog pictures? I was dreaming the other day of putting my film on an IRENE system and being capable to read it using their algorithms. You know, what I mean by that is, image or sound processing by imaging is a very niche market, a very tiny market, to my opinion, it would really make sense if we could just merge the efforts, and merge the knowledge, in order to get something, maybe just one system, or two systems, that, you know, really take advantage of all the beauties that have been developed so far.

Therefore, I'm still dreaming, I'm

dreaming a lot, but sometimes I also get to the point. I'm dreaming of setting up some sort of a center for excellence, for example, in Switzerland, having the two systems, you know, IRENE and VisualAudio sitting together, and somebody, every day, using the two systems, and other people coming in to do some more research and development in order to try to achieve better results.

So I hope that will come true sometimes, but, well, good, that was it. Thank you very much for listening.

PARTICIPANT: Is grooved media still being produced or is there new stuff being produced, I think this is a trivial question, but I assume the body of grooved objects out there is finite, right? It's not growing anymore, is that true?

DR. HABER: I definitely think that it's finite. And certainly, vinyl has had a resurgence. My 15-year-old now wants only to buy vinyl discs, but the target for these technologies, I don't believe, is the new emerging vinyl market. It's the finite historical archive, beginning in the 19th century, and probably ending somewhere in the middle of the 20th century.

PARTICIPANT: Thank you. Stefano, to address the last point that you made, there will be a proposal made, there's a group of scientists, they're having a conference in Padua, Italy on September 17th, and the proposal will address the issue of trying to bring together these scientific communities, and taking advantage of the fact that we don't always have to be in Washington, or Berlin, or San Francisco, but to take advantage of the technology that exists in a virtual sense.

Why can't there be global exchanges of information using something as simple as Skype? And how you integrate the different communities in that manner, for example, in a university environment, you may have an archive, and the archivist has a certain amount of expertise, but the engineering departments, as your example proves, how do you leverage your needs with the expertise that are in the universities?

Or how do you work in New York City in

an archive and find out, gee, I'd like to work with the Swiss archive? And use social media and that technology, I think, could facilitate exactly what you're asking for tremendously.

MR. CAVAGLIERI: Well, thank you very much for that. I think this is a good approach. Yes.

MS. DOCTOR: Hi. I'm Jenny Doctor from the Belfer Audio Archive at Syracuse University. I just want to say, I think you're asking really good questions and I really like the fact that you're saying that we need to use the audio engineer expertise in conjunction with the scientific genius that has been bringing together, sorry, I was using the genius award kind of -- but the point is that there's been a huge amount of approach that has come from the scientific side and I think that there are decades worth of expertise that does come from the audio engineer side that has to do with how the stylus does fit in the groove, whether you do go right down to the bottom of the groove or not, whether some of the noise may be coming from that issue, and whether maybe the edges -- I don't know enough, I'm not a technician, but I wonder if some of the noise problem is because it is targeting the bottom of the groove and that isn't a problem.

So I hand that over to some of the audio engineers who can answer that more, but I know that when we were talking to the engineer at Syracuse University who was doing the laser project that Carl mentioned earlier, and he, for the first time, ten years later, talked to our audio engineer about why that project failed. He hadn't ever realized that he should not be directing the laser at the bottom of the groove. He should have actually been directing it to the wall higher up.

So that was a huge revelation to him, but he hadn't bothered to talk to the audio engineer to find that out. So I don't know about why some of the noise is coming in, but I do wonder if maybe the edges, if going right to the bottom of the groove, if it was slightly higher up, I don't know if you would get better results.

MR. CAVAGLIERI: Well, in our case it's

a little bit different because we are not just measuring what happens at the bottom of the groove. As I mentioned, we are trying to figure out or to set an average between different reference points, but the different reference points are either the top or the bottom, never in the middle.

While, when you use a stylus, you read up on the groove. It's never going to be on the top or on the bottom. That's probably one of the major issues he had.

DR. HABER: So, Stefano, you put up a bunch of questions just before your last slide and I was thinking maybe we could try to address some of those questions since they were questions. Is that okay? Is there some way we could put them back up? Okay. So one of your questions was that the current image sampling frequency is some tens of thousands to hundreds of thousands of samples per ring, while standard audio sampling frequencies are 96, 48, whatever, how do the different figures correlate and what is the impact on the audio quality?

So I would like to say I think the answer to that question is very straightforward. You have a disc, let's say it's 78 rpm, that means it takes some 0.7 seconds, approximately, to go around once. If you take a certain number of optical samples around the disc, by a simple ratio, involving that 0.7, it translates into the equivalent audio sampling.

So I happen to remember in my head that for a 78 rpm, if you take 80,000 optical samples, that is 104 kilohertz. So every single number exactly correlates to an equivalent audio sampling. So 96 kilohertz would correspond to something a little bit less than 80,000 for 78. If it's a 45, it's going to be a different thing, but there's a completely well-defined, clear correlation.

MR. CAVAGLIERI: Yes, but my question is maybe more to him, is, why, while reading the image, these two numbers do not match? I mean, you know, the producing of the, let's say, images, image samples, or just straight lines, is not the same number as the final number of audio samples. I mean, for one second, for example, of audio, with the VisualAudio system, you will not take 96,000 pictures. You will take 120,000 or 80,000.

DR. HABER: Right. But that's because the record is, for example, at 78. Suppose the record was at 60 rpm, okay, then it would be one second per revolution, and then to get 96 kilohertz audio sampling, you would take 96,000 pictures.

MR. CAVAGLIERI: Yes.

DR. HABER: But because you're at 78, it's just a different number.

MR. CAVAGLIERI: No, but if you change the rotation speed, it doesn't matter, because the rotation speed of the scanner is changing in our case.

DR. HABER: It still just matters when you trigger the camera.

DR. JOHNSEN: Yes, I believe it related to the problem of sampling rate conversion. When you have digital sound, when you change a sampling ratio, you have to do a mathematical operation. And where it was done first was when converting 44.1

kilohertz sound record sampling frequency to 48 kilohertz sampling frequency.

And the example from going from 44.1 to 48 is that, if you do it professionally, you can't hear any difference of the record at 44.1 compared to 48 kilohertz sampling frequency. And as long as it's the right way, that you have no problem with sampling.

And in my opinion, that is also the case when you look at VisualAudio system or IRENE. We have a sample rate conversion and sampling rate conversion is nearly perfect. And you can get as perfect as you want, simply, you have to increase the length of your digital filter that does the sampling rate conversion.

DR. HABER: You can run any of these systems if you want exactly at the sampling rate so you don't have to do any conversion. You could choose to digitize 96,000, native, or you could choose, for various historical purposes, IRENE runs at 104 kilohertz, and then if we want, we can down-sample to 96 or to 48, or 41, but you could set it up to digitize exactly at the audio if you want to. I'm sure you could do that.

MR. CAVAGLIERI: I wish we could with VisualAudio.

DR. HABER: Okay. But there's no fundamental reason why you can't do that.

DR. JOHNSEN: We have phase-locked loop and that is making it difficult.

DR. HABER: But you could do it. Okay. So let's go to the second question. Okay. So on an image, the radial position of the groove can be estimated by detecting the edges of the groove on some number of points. On the turntable, it is determined mechanically by the pivot and the inertia of the arm as well as the stylus tracking, so where's the question?

MR. CAVAGLIERI: It's not really a question. It's just, you know, there are two different ways to determine where your groove is and how it moves from one way to the other. The point is, using the optical systems, you have no, let's say, mechanical impact that is, on a standard turntable, designed to produce certain results.

You remove that while reading the image.

DR. HABER: So I think what you're saying is that the mechanical dynamical response, I use the word dynamics in a physics/engineering context, not in the way audio may be using a dynamic range, in a physics dynamical sense, you don't have the compliance of the mechanical components, the tone arm, the stylus, the mass, the natural frequencies, so these are all parts of a physical system, which together, conspire, hopefully, to give you good results.

There are many examples in the literature of problems associated, resonances, and distortions, that people have worked for years and years in perfecting mechanical playback systems to, you know, minimize. Those effects are not present in an optical scanning and VisualAudio, or IRENE. All you=re doing is measuring the shape of the groove.

Mathematically, the sound should just be the derivative of that shape. And when you take the derivative of that data, it is an absolutely mathematically flat in frequency. And I think when you were using the term linear transfer, you mean a flat transfer. It's completely flat.

In an ideal world, the cartridge would simply take that derivative by the fact that it's the velocity between the magnet and the coil inside creates the voltage that gets measured, and that's what you get. If you want to add the flavor or the aesthetic, if you like, of a mechanical system, that's called physical modeling and you can develop that, but it's not there. It's not there.

Then it becomes a question of aesthetics, right? Do you want that ambiance? Because you're not going to have that. It's just going to measure the information that's on the disc or the cylinder.

MR. CAVAGLIERI: Yes, I understand that, perfectly understand that, and this is actually my problem. I mean, because also when the groove has been cut into the disc, it has been done with mechanical means. It hasn't been in an ideal world.

DR. HABER: But a difficult mechanical system cut it than you're using to play it back. And the mechanical system in Indiana, and the mechanical system in California, and the one in Switzerland, they could all be different, so in some sense, some people might argue that it's good that you've removed one set of variables from the playback process, which is the mechanical response.

Now, you raised the question about noise and you very dramatically showed these high noise levels that you're seeing, so I was thinking I would like to address that also because I think it's part of your question. When people use the term noise, it refers to a whole variety of things.

So if you have sharp defects, click, click, click, okay, people call that noise. If you have a very broadband, like the hiss that you so strongly showed, okay, they also refer to that as noise, but these are very different things and they all get stuck in the same category and called noise.

What you were dominantly experiencing

or sharing with us in the thing was this very, very broadband noise, and it was very clear when you showed the noise spectrum that at some point, around a kilohertz or something, it started to rise and it rose, sort of, in a straight line, and it was much less prominent when you, for example, in the test record, played it back with a stylus than when you did with the VisualAudio system, okay?

That noise is due to just a structure that exists in a random way throughout the data, okay? And it can be due to, perhaps, the film grain, it can be due to the roughness of the surface, it can be due to the region of the groove that IRENE or VisualAudio resolves, but it's a general structure of the surface as rendered through the imaging which gets worse and worse as you go to higher and higher frequencies because you're taking a derivative, a derivative is a high-pass filter, so you're pushing up those high frequencies, which you should be doing, okay?

But you should apply the RIAA curve or whatever the corresponding curve is for the older materials, which is, in fact, designed as part of the standard process to partially roll off those high frequencies. It does it for the stylus and it should be applied in these visual as well. I don't know if you were applying it. It didn't look like you were applying it, because I would have seen it turn over, but it definitely should be applied because it's part of the way the record was recorded.

MR. CAVAGLIERI: Well, yes, of course. Thank you for this explanation, but it doesn't really answer my question because we are well-aware of the fact that we have to apply compensation curves. The problem is not there because the compensation curve in this case removes everything. It removes the signal as well.

And, you know, the signal-to-noise ratio, especially in the higher frequencies, is that low that you can't remove the noise without removing what's embedded into the noise.

DR. HABER: Sure. Once it's below the noise, if you put a curve on it, you're not getting it back.

MR. CAVAGLIERI: But the two noise floors I showed were obtained using the same equalization curve, you know, the one from the turntable and the other one from VisualAudio.

DR. HABER: Okay.

MR. CAVAGLIERI: So it was already in.

DR. HABER: Okay.

MR. ALYEA: Okay. We should let people get to lunch because we need to be back here at 1:30.

MR. CAVAGLIERI: Thank you.

(Whereupon, the foregoing matter went off the record at 12:21 p.m. and went back on the record at 1:35 p.m.)

MR. ALYEA: Also, at the end of the sessions today, there's an establishment on Independence Avenue called the Hawk 'n' Dove, that if people are interested in meeting outside of the walls of the Library of Congress, maybe be a little looser or something, I don't know what exactly, but, you know, say hello and be casual, the Hawk 'n' Dove will do that, so I guess we'll probably be out of here, what does it say on the schedule, maybe 5:00, 5:30, so right after that, and then people can go off to dinner.

So we could go directly or a little bit after that, so 5:45, let's say. Split the difference. So we should have the next set of speakers come up, John McBride, Carl Haber, and Stig Molneryd. So, Carl and Stig, could you also come up and sit up here? We're going to try to get everyone cycled through a little bit faster and they'll all be ready for questions at the end.

So we have John McBride. He's going to be our next speaker.

DR. MCBRIDE: Okay. Good afternoon, everybody. The first session after lunch is always the worst one to have, especially if you had a large lunch and you're beginning to fall asleep, so I'm going to try and livening this up a bit with, hopefully, some interesting material for you.

What I should say here as an introduction, I'm an engineer. I've only dabbled in this area for around five years and the project I'm going to talk to you about here actually concluded around 2010. I'm going to give you a bit of an update as we go through, but it was just a piece of work I did from my main job, so I'm a research scientist.

I work primarily on structured surfaces, on surfaces, and of course, recordings are structured surfaces, and structured surfaces are really a big area of research at the moment for all kinds of things, flight, ships in water, everything, so this is just kind of a side road for me, but I hope you find some of the work we did interesting.

There's a bit of a history behind this. Just get my notes here. Can I have the first slide? So I'm just going to give you an introduction, and overview, and probably just focus on the case studies. Actually, Carl here has done me a great favor, and Ottar as well, you covered most of the technical aspects, so I'm going to kind of leave that.

I've got a lot of slides here, which, I didn't really know what I was going to say today, so I'm probably going to jump quite a few of these slides and just go into the interesting stuff. I'll probably focus on some of the recordings we looked at, particularly the Queen Victoria, the tinfoil recording, and maybe the Berliner disc as well, but we'll come to those shortly.

Before that, I think the next one is a video, which I hope is new to you. One of my PhD students found this on YouTube. I hope you like it.

(Video played)

DR. MCBRIDE: So some of you have seen that, I guess, but if you haven't seen it, I think it's on YouTube and it's worth it. That's the way to damage a cylinder recording, I guess. Here's another one, which I'm not going to play, but Carl would probably be interested in this. This is a Radio 4 Today. It's the most serious radio program in the U.K.

It's early in the morning, you get presidents, prime ministers, it's really serious, and this lady here is one of the serious people who present the program. I think about five years ago she was giving a talk on, you know, it was a very small slot on one of the early sound recordings, and I think it was one of the optical ones, and she gave the one-minute slot on this piece of news that somebody had managed to decode this recording from 1850, or something, and the guy in the headphones said, it sound like a bee in a bottle, and she burst into laughter.

And she then could not stop laughing. She had a giggle fit, which actually went on to the next part of the program, which was an obituary, so she's reading an obituary of some famous person who passed away, laughing as a result of this. So if you want to find it, Carl, you'd probably be interested in hearing that.

DR. HABER: I heard it.

DR. MCBRIDE: You heard it. It's good. Okay. I'll come back to more technical subjects. So probably, I can miss this. This is really just saying that what we're trying to do with this piece of research is not real-time playback. This slide was for a general audience, so I'm going to skip most of this. It really describes what Carl said in the first presentation. Let's just move on.

Okay. I'll jump this one. There is a Web site, which we still have active, actually, which, if you get the slides later on, we should be able to logon to this, and all of the material that I'm going to show you is actually on this Web site, so you can look at this on your own leisure.

I think the key thing is archivesound, one word, and if you Google that, you should come to this Web site. Okay. Let's move on. So as part of the project, we published quite a few publications. These are primarily in the scientific literature domain, so those of you who may be engineers or scientists, a lot of this work is easily available. It's downloadable. So a lot of the technical content is published.

I won't bore you with the detail of these, but maybe make reference to a couple of these papers as we go on. You've seen this. This is a measurement method. Carl explained this earlier this morning. I'll probably say a few words about some of our systems.

This particular system is very similar, I guess, to what you call IRENE, which is a white light sensor with a cylinder. This has got a cylinder on here. I can't see it very well on the screen here. I hope you can see it on your screens. There's a cylinder here being scanned.

And these axes are really important, so X is a long cylinder, obviously, theta is the rotation of the cylinder, and Z is the height position, so we're actually measuring, as we scan along the length of the cylinder here, we're measuring the Z height as a function of X, and then you rotate, and then you do the next scan.

And from that, you essentially are creating a three-dimensional representation of the object. Let's just move on. Another system, I'll have a movie of this system later, this is a flatbed scanner. It's an air bearing, which is a high-end piece of equipment for -- actually manufacturing silicon wafers, but what we do with this, as you can see on here, probably this is a 78, I guess, with a laser, and again, you've got the rotation, X and Y, you've got Cartesian coordinates with a Z height.

Okay. I'm just going to pick up on a few of the aspects which we've kind of discussed a few times from the presentations this morning. Something about resolution, and maybe resolution is the main one we should focus on, because this is taken from one of our papers, the reference is there. I hope you find this is interesting, this graph is essentially the physical displacement of a cylinder's surface measured in nanometers again the frequency response of a spoken word.

Now, if you talk to people who are experts on spoken words, the harmonics -- so here, you've got the main sound here, which is around 600 hertz, these are harmonics, the peaks that occur afterwards are the harmonics, and here you see a harmonic around, what, 3 kilohertz, and if you look at the amplitude of this harmonic, it's around 50 nanometers.

This is a recording from early 1900s,

it was a cylinder recording, and what we determined from this is that, actually, if you're going to measure these surfaces and get a good audio record from the surface, you really need to be measuring to this type of resolution, so we published this, I think, 500 nanometers was the value that we suggested from this work.

It's to do with the spoken word and the way in which the ear interprets the spoken word that these high frequencies, these harmonics, are really important. Let's move on. One of the other issues we haven't seen mentioned this morning is the optical sensors that we use are not perfect. In fact, they have some major limitations and one of the limitations is demonstrated very easily here from a 78 disc where, as we saw again this morning, the side wall here is at 45 degrees.

Actually, it's very difficult to get light back from a surface that steep, and what you see is you get a lot of data, basically, on the side walls. Now, from what Ottar presented in his presentation, this probably isn't that important because, actually, if you can detect the edge, you can get the sound, actually, from these edge points, and also from these points here. We have an example of that later.

This is signal extraction. This is really a whole area in itself. I could talk to you for an hour just on this; how you do this. In the simple sense, it's actually what we've said all along, all the way through this conference, basically, you're getting a three-dimensional representation of the surface, and then the secret is to find the groove, and there are many, many ways of doing that.

And you're really entering the realm of signal processing, and anybody who knows anything about signal processing, it's endless, the possibilities that you can use. The one that we tended to focus on is one particular -- so the dark line on here is a cross-section of a groove structure, and I think you can see, it's damaged.

The bottom of the groove here has quite a strange base. It's not curved. And you can see if you just try to find the minimum point as a representation of the sound on this, the minimum point is actually going to be noise. It's actually a damage to the surface. So actually, what we've been doing here is putting a particular filter to the surface and getting the minimum points of the filter.

Again, there are endless filters that you could use for doing this, but having got -- the key thing is, you have the raw data here, the solid line, as accurately as you can get the data.

Okay. So that's really all I'm going to say about the technical aspects. I can talk, as I said, if you have any questions, I can go into all the details if you want. What I'd like to do is just focus on some of these, which I hope you find interesting. I'm going to jump the first one and not bother with that.

There are lots of cylinders that we scanned as part of this project. They're all available on the Web site, so if you want to listen to them, you can. Just Before the Battle, Mother, is, I think, the first one we did, Carl. That was a collaboration with Carl Haber, with Bill, the cylinder supplied, so we did that, probably, 2005, I guess. Something like that. So there are lots of other examples on here.

This is quite interesting. We saw this partly this morning. We actually found a character in the U.K. who made test cylinders, and this is one of the cylinders that he made for us, so these have fixed frequency sounds embedded in the groove so that we could do a way of testing the quality of our sound reproduction from these scanning methods.

I think you can just about see that. If you look across the cylinder here, you can probably just visually see. I haven't got it very clearly on this monitor, but on here, I think you can just about see the different regions, and then we've scanned the cylinder. And what I'm going to show you, so here's a groove structure from the cylinder.

Let's just jump on and listen to some of the sounds. So I'll just take one example here.

If we look at the 1 kilohertz. I'll play the stylus first, so this is the stylus playback of the cylinder we had made. What do you hear? Can you hear a single frequency? It's modulating. Yes. There's some variation. Let's listen to the optical. Same frequency.

So we've written a paper on this. So one of the publications, I think it's JASA, the Journal of the American Acoustical Society. It's one of the top acoustic journals in the world. There's a paper published on this work. So we created the artifact so that we could test the quality for the methods that we were using.

And actually, from this work, we actually determined some quality parameters that you could use in going back to the measurement and to the way in which we process the surfaces afterwards. I could go through all of these, but let's just jump on.

Okay. So these are frequency analysis diagrams of the same data. If you aren't used to seeing this type of stuff, this is the one that we've just been looking at. I think the 1 kilohertz, around here. So you've got the stylus scan. So this frequency here, where the arrow is pointing, is 1 kilohertz, and then you've got the harmonics, which is probably the sound variation that you were hearing, and if you compare that to the optical, there's more noise on here, you can see slightly more noise, but the primary sinusoidal signal is well represented there.

As I said, you could read the technical papers on these if you want. This is one that is really interesting. I have quite a few slides on this one. So this was scanned, I hope you can imagine, Queen Victoria is quite an important character. It is for the British. And there was a book, I think, written around the recording of her voice.

But at the time, it wasn't sure if it was -- it hadn't been authenticated, and I'm going to show you some evidence later about the authenticity of this. So this was scanned at the Science Museum in London back in 2005. I just really want to focus on -- there's the cylinder. This is a graphophone cylinder. I think you can see three dams, one here, on the left, one in the middle, one on the far right.

And I hope you can see it. From here, I don't have it very clearly on here, but I know the slide pretty well, and you can see the groove structure. And what you can also see is this artifact on the top here. You can guess what that is. That's the stylus. So essentially, what has happened, this artifact has obviously been deemed to have been important. It was in the vault of the science museum in a huge safe.

And I think, probably, at some stage in its history, it'd been played, probably played with the wrong stylus with the wrong force, and what this has done is, essentially, damaged the bottom of the groove here. That was our interpretation of that.

Let's see the consequence of that. Well, this is the consequence right here, so this is a cross-section from the data. This would be the top surface here of the cylinder, these are the grooves, and you can see, I think what you can see here, this may have been done in the distant past, in fact, because it looks like a piece of wire has been used to try to play this.

And I think, Bill, you were saying that piano wire was one of the early stylus that you -- steel wire? Yes. And you can see what this is. You see this repeated shape on the bottom of the groove is essentially like a piece of wire that's been cut. And what that's done is, essentially, removed, of course, some of the key content of the material here.

But that's one of the great things about doing three-dimensional scanning. If you look at the -- there's a virtual stylus. That's what's been removed. But if you look to the left and right of that, let's just go back, okay, this should be, okay, here you go. So you've got the left-hand side, the right-hand side, stylus in the middle, but you've got this whole region to the left and right of the virtual stylus, which is essentially where the sound would have been carried.

And because this is all in the digital

domain, we scanned the surface, we can then, actually, decide where we run the tracking for this to get the sound from the groove. It was quite damage, so we're not going to hear anything, but I'll play you something and you can judge. So this is pretty much what we saw this morning as well. If you want a binary indication of the surface, so the red is the left side, the right is the blue side, and there's the green, is the track of the stylus, and then all we're going to do is we're going to move that to left and just track to the left.

Okay. So let's listen to something. I'll actually go straight to the end. I won't go into the complexities of this, but I've kind of kept this secret. I've played it a few times in the U.K. at sound archive conferences, but we've never released this, and it's certainly never been played outside of the U.K. Listen to it first and then we'll go to the words.

Okay. So what we've tried to do here is just interpret what we think is being said. So greetings is quite clearly stated. We can hear a spoken word, but we can't hear what's being said, "the answer" is very clear, and it's said with a very high-class British accent. "Can be" is also clear.

This, we're unsure of, "Lord Granville", I think, was the prime minister. I have to go look in my history books. "Absolutely", is very clear and "has never forgotten" is also clear. Let's play it again. Okay. There was some uncertainty about who was speaking. And actually, the book that was published on this, left the question open.

But then, since the book was published, and I think that should be my next slide, yes, here. This is a letter I wrote to the Royal Archives, this is April of 2007, and I've just highlighted part of this text here, which is the next bit, so the first bit in yellow is referring to Paul Tritton's book.

And it's talking about a letter from the Queen's private secretary, Henry Ponsonby's wife, Mary, which came into the archive after the book had been published, and essentially, this is the key bit on the bottom here. "HM", Her Majesty, "spoke into it." So this was on the letter from Henry Ponsonby to his wife. So we know that she spoke into this device.

And I think if, I may get the dates wrong, but sometime -- here we are, 29th of August 1888, so we know that she spoke into this. And actually, she said didn't want to -- I think what she said was, she didn't want to -- yes, here we go. "But we told Mr. Morse he must not go around the country producing the Queen's words." It's a very early piece of history on sound recordings.

I think actually, the machine was a graphophone. It was one of the artifacts we saw in the museum yesterday at the Smithsonian. It was actually taken to the U.K. from the U.S., and then taken to the Queen so that I think her and her husband were keen advocates of new technology. Okay. So a bit of history.

This is another particular piece of history, if you're Welsh, or if you've got Welsh ancestry. I don't know who this person was. I have to go back and remind myself, but this came from the Welsh National Archives. And essentially, this was Evan Roberts, who was a preacher, 1905, but the artifact had been broken.

So it wasn't the one that was in the movie that we saw in the beginning, but this is the casing that held the artifacts of the broken cylinder, and this was held in the museum in Cardiff, I guess, in Wales, and it was sent to a dentistry expert in Los Angeles who did restorative dentistry, yes, and this is what came back.

This is before we got involved, so this is actually what we were presented with. Again, I can't see this very clear on here, but I know the image pretty well. I think what you can see is areas of filling, essentially, and I think you can see the damage and I think you can see regions of growth, mold growth, around the cylinder.

But again, with this type of artifact, we're able to, basically -- ah, here. That's a better picture. Were basically able to, like we saw with the previous presentations this morning, we've got part of the groove structure, and then we can do interpretation, so there's no groove here, but we can then add the time period in as a blank and actually get the sound.

And here's a good example. So here you can see, let's get the trajectories, so the blue line would be the conventional line, and then you can see, this is the broken surface here, and then we've just changed the trajectory of the stylus. You can see where these things -- I can't see it very well on here, but I think you can. Hopefully you can see, to the left, you can see the structures basically follow, and the green line becomes the virtual stylus.

So this is a great advantage with this technology that you can use by software and techniques, basically, to interpret the data. Let's just play one of these. This is optical. So I think I recognize that. It's a song and I think which used to be like a church song, For Those in Peril on the Sea, so that's recognizable. This is the stylus.

So at the time we were doing this project, it was quite frustrating because we developed the technology, but everybody who came to us came to us with an impossible sample to measure, so this was an example of something that was really difficult, like the Queen Victoria recording and this one.

And then we were presented tinfoil recordings, came our way. This is the U.K.'s oldest recording. So it's an Edison tinfoil. We were presented with this to see if we could scan this. Because we did that, actually, the sound from this, I think I have it here, but because we've done this piece of work, we were then contacted by the Norwegian Archive. I don't read Norwegian. Can you read it?

It's 1879. I can read that. Okay. We have the date up here, so 1879. This was quite interesting. I think I have this one. First, this is the system we used. Sorry. This is the U.K. one. This is the Edison sample from the British Library.

I think this is on an exhibition at the British Library in London, and actually, the sound from this is also at the British Library.

Okay. So you can see how that worked. This piece of equipment is an air bearing. Very expensive piece of equipment, but it's basically used for silicon manufacturing. It gives you a very flat surface, very high speed, and it's that type of technology we were using for doing these surfaces. Oh, this is one I actually was presented with in the U.S.

This is from, I'm not sure that we can see it very clearly on here. We didn't do this one because it was so badly damaged. This is held in a museum, I think, in New York, and it's very damaged. I don't know if you've seen this, but they were asking -- you did. Okay. Did you manage to get the sound?

DR. HABER: Yes.

DR. MCBRIDE: Okay. That's great. Because we looked at and thought, gosh, this looks too bad. So this is the Edison recording, 1877, badly folded and ripped. Again, the sound is on the Web site. This is the Norwegian one. So I think it's somebody playing a Norwegian horn. Yes? Is that right? So this was presented at an archive conference in Oslo at the end of this project and I think it's in the science museum in Oslo with the sound.

But this particular artifact was in very good condition compared to the British one. Okay. I'm just going to come towards the end here. I'm going to jump the 78 disc and just go on to the Berliner. So the EMI archive is right next to Heathrow Airport in London, and they have a museum, which is not open to the public, and it has some of the most incredible artifacts I think I've ever seen, just around this room, and some of you may have seen it. It's just incredible.

Everywhere you look there was some incredible piece of history from sound recording, and it was just in this vault, and they gave me this to play with, which is a Berliner. I think this is one of the earliest examples they had. I'm not quite sure of the history of this. I was looking online just to remind myself where this came from, but this was a metallic master used for printing, probably, on rubber.

We actually scanned it, but we never actually got around to finishing this piece of work off, so we have the surface and we just haven't really got around to finishing it, but I just wanted to show that to you as an example. And I'm just going to finish with, I should say thank you, actually, to all of the people who have helped me with this.

There were too many people to mention, but one of my PhD students, Anthony, who did his PhD in this, went on to work in the film industry, so he did a PhD with me on early sound recordings in engineering. He now works for a software company in London doing CGI. And he sent this to me, this movie, which I like.

Okay. So what he's done is taken raw data from a scanned surface, put it into CGI software, and it recreated the surface. So you could fly through this surface, you could do whatever you want on a movie with this. You could actually follow the grooves or whatever, so he just sent this to me as an example.

Just actually, by coincidence, I said, oh, I've got to go present something in the U.S., and he sent me this. I thought you'd like to see it. I really like it. What it says is something about the future of what we could do with some of this scanned technology. This is maybe trivial, but it could be something that would be good in a museum as a demonstration of what these things are. And this data is raw data. He's just created it as a movie as part of a CGI. I can't remember the name of the company, but it's quite a well-known CGI company in London. He's probably working on the next Batman movie, or James Bond, or something.

I think that's it. I hope you've enjoyed it. Thank you very much.

MR. ALYEA: So now we have Stig Molneryd. Did I get close?

MR. MOLNERYD: Yes. It's nice to be here at the conference. I want to thank the organizer for this conference at the Library of Congress. It has been very interesting days for me and to participate in the conference, and I also made some other visits here in Washington, D.C. I've been up in Boston and Andover seeing the IRENE technical part of it. It was very interesting.

I've also been in the National Archive and seen how they have been working there. And yesterday in the Culpeper, it was very interesting to be there also, and heard the discussions about all the preservation problem they have, as we also have in our archive.

And there are many things to go on with together, I think, to manage all these problems. And this is the new possibility to make imaging of analog mechanical records and we have just started, just started, we are just finish it in the end of August this year. And I will have a short presentation of how we have worked and what the results so far have been done.

Yes. And I'm going to do brief presentations of the innovations and development project started at the National Library of Sweden called KB, it's Kungliga biblioteket. The project is on innovation procurement, PCP, entitled, Development of Contactless Playback of Analog Records to Digital Sound Files.

And about myself and my background, I've been working at many different types of media and carriers, film, song, video, and I started in the early '70s. And I have some experience about these fields and preservation of these things.

And I've also been interested in technical developments in all different fields, and this project is one of these. And now, yes, here we are. We have, in the archive in KB, and I will show you how it all started. We have a total amount of about 146,000 Swedish records and of these, there are 80,000 78s and 60,000 vinyl records, EP, single, LPs, and so on, and about 6000 lacquer discs.

We are doing some cooperation with Swedish radio and they have 285,000 records, and they also have foreign records. And depending on our expansive archive of records, we want to investigate whether it's possible to find ways to transfer these faster than real time and examine the sound quality that can be achieved.

It all started when we were contacted with a company that showed a demonstration of optical playback of a 78 record. It was surprisingly good quality. And the technical development for image scanning used by this company was taken directly from the shelf and had not been optimized for the applications we requested.

We had discussions with the company and described our needs to be able to transfer the disc faster than real time and at the best possible sound quality. They also indicated that image scanning operations and transfers could be faster than normal playback and the possibilities to edit and optically clean the discs.

After that, we were looking for sponsors to investigate if it is possible or the technology could have what we asked for. With this material and proposals for the development project, we ended up finally to the innovation agency in Sweden, VINNOVA, and we made the application with a project description of VINNOVA, and got it granted as innovation procurement, PCP.

We had no previous experience in innovation procurement, and it has been very interesting to run the project with the development process that applies to this type of pre-commercial procurement. During the whole project, there had been a form of cooperation and dialog, but also involved clear boundary issues of secrecy for the companies.

These connections have created a good and challenging development and exchange of information. And when we started and we got this information, you can see the picture. It was this type of just turntable and I have the lights, and the camera, perhaps I have it better on the other pictures. It's very dark here.

And it listened like this. And with this, we went on and we will make a short description how the PCP project have been working on and the result we want to achieve that we were running at

KB. It's an innovation project with funding from VINNOVA and also from KB that started the 1st of January 2014, and will be ended in the last August this year.

VINNOVA is a governmental institution in Sweden that provides financial resources and opportunities for companies and governmental institutions to apply for funds with development of various innovations and investigation project. And I can say today, in EU, it's focused on innovation development, and has dedicated enormous financial resources that can be applied for development and innovation proposals today.

It's a special program they have announced. We run the development project ourselves at KB, but have, in our reference, an expert group participants from the development department at the Swedish radio and also external experts in optical scanning and image recognition.

It is an innovation project, which means that the result will be presented as a description of how far the technology has come today and can achieve with picture scanning and with development in image processing. The project will demonstrate the results and the possibilities that exist to get audio files from analog records.

Our goal is to get results which show faster than real time with best possible sound quality. And so we started this procurement and we received a total of eight proposals from different companies. They had new, earlier professional experience in audio recordings and were more oriented in the industry in picture recognitions.

In Sweden, typically, an example of sorting and selecting knot-free woods, planks, metal industry, productions, detections, robotics selections, failure in productions, and so on. And we selected all these five companies with innovation solutions that processed to phase 1.

In the continuing work in phase 1 includes producing a report on a preliminary study and detailed description of its prototype solutions. After examination of these phase 1 reports and the associated round of interviews by these companies, we made a second new selection and we selected three companies, which now are participating in phase 2 of development and production of prototypes.

They have to make and produce audio files faster than real time from 78s and vinyl records. In phase 2, includes a number of goals that should be reported and be checked before the final report. And we have to note something about the result and we have give them these two records, test records and reference records, and it's important that we have good reference and test records to compare with normal turntable playing or with ELP laser turntable.

They had got these records to work with and compare the bandwidth and other parameters. It's the way for us to just compare the result we get from these companies. The three companies now participating in phase 2 have different proposed solutions and prototypes that make it very interesting to see how the results may differ.

They're working with 2-D and 3-D

scanning methods, laser or LED lights, and other tools that they have in their prospect. And so far, I'm not an expert in picture recognition or image scanning, and so on, and they have to proof and show the result for us. And we can only see the potential for the results when they have delivered the final report.

And I have understand so far it's much about the lighting, high light intensity and the brightness, and the problem with unfocus and blur, and also, the pixel size in the cameras, and so on, have limitations for this work with imaging. And one company have used laser and there was a problem with speckles, that I have understand it's a problem, also, about the quality.

And we go on and there are three companies, and the first one here is MBV Systems AB in Sweden, and they use a 3-D laser triangulation, and they use also, what they call Scheimpflug techniques, and it's for 78 records. And if you have questions about what the technique is, I can't answer it. You'll have to contact this company.

Perhaps in the end of this project, I can more about how the results will be describe and so on, but this is the hardware, you can see here, and they have a laser and they have done these things. And I must also say that these sound recordings and so on I got is from some months ago, and I can say that it's a rather short barrier for them to manage and give us. It's almost half a year.

And I think they were very glad when they -- and they have built-up this hardware and get the sound from it, and you may listen to what it looks like and hear. This is from the MBV Systems. Very short, but you hear something. And the next one, the second, is InVite Vision AB, they use 2-D, and they have mostly worked with LPs, vinyl discs, so far, and this is the hardware, you can see.

They also have -- I'll go on here. InVite Vision is another company working with through the years, and have focused on primarily vinyl records. Pictures of the whole equipment with label camera. This transmitter will display LED light camera and servo motor adjusting the distance. The disc is moving with an air bearing, compressed air, and rotate. About 20 grooves in the detection picture.

I forgot to say something about this I have right down here. They are working with the 78 records and their plan is to make it ten times faster than real time and catch 20 to 100 grooves in a picture. That was what they have told me so far. And third one then is -- we can listen to the InVite Vision also.

This was vinyl disc. And I think they have the same recordings. And this, what we were listening to, was in a vinyl recording; vinyl disc. And then third one then is Joroma AB, and they called this system Optofon. Joroma is another of the companies working with 2-D. Reported so far about hardware competence and improved algorithm for pictures of the grooves, et cetera.

They are building the prototypes, and are doing tests, and are planning to do tests with broken and cracked discs. The plan is to make 20 times faster than real time and to catch more than

40 grooves in a picture. Joroma AB is a Swedish company with experience in developing automatic inspection system for industrial applications. Theoretically, able to capture all grooves in two revolutions.

New processing tools can be reapplied on already scanned records. Individual grooves are extracted and tracked, as you can see in the picture, and the image processing, sound is extracted using machine learning, I'm not sure about it, probabilistic methods to avoiding defects in grooves. It's only something I'm not common with.

And then sound processing, sound can be post-processed used in signal processing. Raw data is always preserved; untouched. Can be useful to apply noise filtering. And they can be contacted. The man who has been working with this is Josef Grahn in Sweden in this company. And I have no audio to replay from them so far.

And the next one is then to talk about is the, as we call, same question, different solutions technical limitations. The companies shall submit the final report on the 24th of August with presentations on results achieved and suggest future technical solutions that they have developed.

After the project has been finished, we are going to evaluate the final result to see if and how we will proceed in our continuing work with optical playback and digitalizations. Depending on the results achieved of sound quality and how quickly and efficiently they may transfer the records, we evaluate how realistic the applications and automatically those digitalizations could be designed.

It should be clear from the companies' final reports that their optical playbacks and technical solutions have the potential to achieve the quality we demand. There can be results that we get from the optical playback of 78s that will achieve the quality we demand, but not for transferring the disc and so on.

If the results are not going to be what we want and be good enough, should the reasons be described what it depends on. It may be that new technical possibilities of optical scanning technology will be further developed in the near future and provide new solutions. It can then, perhaps, create opportunities for optical transmissions and image processing, and provide the results we are demanding of full quality during playbacks, even faster than real time.

And this is what we hope and perhaps we, in the future, have some robot to make this to work. And we are planning to work, depending on the results obtained and the potential that may exist, we plan to organize a workshop during the late autumn this year, where presentations of the results that emerge from the participant companies.

And all these information we have heard here from IRENE project and so on, we will take with us and add to all the result we get from our project, and see what we can do with all these results for the future, and we never know. We hope it will be what we want.

And I'm eager to -- I don't know the

results from the companies. I'm eager to see what they can present in the end of August this year, and see what we can go on. And that's what we are working with and so thank you very much.

MR. ALYEA: Now we have Carl Haber again.

DR. HABER: Okay. I said good morning, now I say good afternoon. So this presentation will focus more specifically on the IRENE technology, which is a high resolution direct imaging, meaning, we image, directly, the media. So IRENE is a set of tools, it's not one thing, for non-invasive preservation, restoration, and transfer of mechanical sound carriers.

It employs direct imaging of the surface in both two dimensions and three dimensions. It consists of hardware, which you can call the scanner, and a variety of software. There's software that controls the collection of data and the control of the scanner itself. And then there's software that is used to analyze the data.

There are five installations of this

equipment in one form or another, in Berkeley, there are two at the Library of Congress, there's a system at the Northeast Document Conservation Center, and at the Roja Muthiah Research Library in Chennai, India.

The IRENE project is a collaboration between Lawrence Berkeley Lab, the Library of Congress, University of Applied Sciences in Fribourg, Professor Johnsen and the Smithsonian, and something like 40 students have participated in aspects of this work, and then all the different agencies and mostly public, and some foundations, that have supported the project are indicated here as well.

So this is a picture of the IRENE system, which is actually at NEDCC now. So it's kind of what you would expect. You've seen similar pictures in the other presentations. There's a vibration stabilization table, which is on air, there's, essentially, an arch that contains both the confocal probe for doing the 3-D imaging, and a camera for doing micro-photography for the 2-D imaging, there's a turntable that holds disc media, there's a mandrel that holds cylindrical media, both of these probes can be moved around under the control of a computer.

There's some lasers that are used to control focus and there's other stages that are used to move the cameras and the probes up and down. This is a picture that's supposed to evoke the three-dimensional data where you get the full topology of the surface and this is a picture that should invoke the two-dimensional imaging where you're just looking at the light that's directly reflected back at 90 degrees.

There's various control computers and so forth, illuminators, that are part of the infrastructure. So this system is quite general and it images a whole variety of media and the characteristics of that media affect the quality of the sound that you get. And there's no one number, there's no one result, that describes how this system works because there's so many different variations that occur.

Just to group a couple of prime things here to point out some of the differences that you have heard this morning and you'll also hear when I play a few samples later. So lacquer discs have a very, very fine-grained material that makes up the surface, and when you image them, so this is a blown-up region of the groove, the appearance is very smooth.

Shellac is made of, you know, crushed beetles, and a bunch of other things, it's a highly composite material which has a lot of texture, and particularly, as it gets worn down, different elements of the matrix wear at a different rate. So you have, instead, when you image shellac material, you see already a much more irregular structure.

And when this gets turned into audio, there's going to be more hiss, the broadband noise that you head some examples of earlier, in shellac materials than lacquer materials. And it goes back to, essentially, the quality of the material and how that represents itself in an image.

When you're in the 3-D domain, now, this is not a photograph anymore, I remind you from what we saw this morning, it's a depth image, so dark is deep. This is a curve that shows how the surface is modulating. Here you've got many, many points that describe this surface. It's not simply an edge transition anymore, and so you're able to do a lot of averaging when you have that much data, and that, itself, has an effect on the noise.

And you heard some examples of cylinder data this morning when we looked at the Rakoczy March in the context of talking about frequency. This is another type of three-dimensional imaging. This is an aluminum transcription disc, so it's a groove that's actually a lateral groove, but it's embossed. So you have soft aluminum and the cutting tool doesn't remove material, it just moves it.

And so what the groove looks like is, it's a depression, like a groove, but it's always got these horns on it. Okay. The horns are the material that's been pushed up out of the aluminum and now sit up on the side. If you shovel snow, you know, it's the stuff that you pile up along the sides.

So the audio is stored down in here where you have many points that you can use to characterize the groove and then that has an effect on the sound. There's a lot of software that's come about, both control, data acquisition, and analysis. So just to give you an example of what the operator might encounter, this is the analysis software that's used to process 3-D data, so it's one of the software tools.

So what you've got here is a panel with a bunch of controls, and this is a summary of the data that was taken from a cylinder, so imagine the cylinder, you slice it and you unroll it, a cylinder becomes a rectangle when you do that, so these stripes indicate the groove and the top of here would go down, and then it would join again, so this is a mapping of a cylinder to a rectangle.

If you zoom in, and you can see this, and you can see a lot more detail of the surface, and you can see all this mold damage. I had talked about that this morning when we looked at the Ishi cylinder, and so this gives the operator a much, kind of, more detailed understanding of what's going on.

And here, along this line, you see the actual up and down profile that, roughly speaking, tells you where that stylus would be moving and how it would be moving. So there's many, many options that this gives you and I realize that somebody this morning was saying, there are too many options. And in fact, that is an issue, and what you would imagine in some context is that this thing converges on a much simpler set of presets, and controls, and even becomes an application that can be distributed and used along with the image data that, clearly, people said that they would like to have.

This is another view of another aspect of the software. This is a chart which explains all the software tools that were created with the design and specifications of the Library of Congress to enable them to operate the IRENE machine that some of you saw when you went down to Culpeper, in a production workflow sort of environment.

So the red flowchart is sort of what the operator would be doing when they're scanning a disc, and the grey stuff here is the analysis of the data that's being created that's happening automatically in the background. So here's some database and the operator would enter some information, and some information would then be pulled in from a database through this front panel, and then there's some automatic, kind of, calibration procedures that go on that quickly figure out where the disc begins and ends, and exactly where to focus to start off the acquisition.

IRENE, as I said, uses direct imaging and so focus control and depth of field are technical specifications that are much more detailed and important to the way it functions. And when you have a lot of depth of field as you might have, for example, in visual audio. So it automatically finds the record, this is sort of the beginning and the end, then it figures out what the focus is, and then it, basically, gives the operator an image quality feedback, and then the person executes the scan, the data gets written into a file, and there's a program that's sitting there looking for new data.

When the new data comes, it analyzes it by pulling in all the images, edge detecting, tracking the position of the groove, cutting out the noise, if necessary, and creating an audio file.

So this is, I guess, with some, probably, improvements, because this is an older slide, similar to what NEDCC is using. This kind of tool set. Okay. So let's listen to an example. So this is a laterally cut shellac disc, circa 1930, which was in pretty good -- no, I would say moderate condition, so it's imaged using the two-dimensional aspect of IRENE.

So you're basically ending up with pictures like this and those pictures have high contrast. They're kind of like the pictures you saw in VisualAudio, but they're directly imaged. There's no film, intermediate film, step, and edge detection is applied, in this case, just to the bottom of the groove, not to the top of the groove, which we view as more worn and potentially more of a source of noise.

So if you play this record with a stylus, it sounds like this. So I don't know if you can notice with the sound system, but there's, obviously, the audio and then there's some crackle, which is kind of a, not a broadband hiss, but a kind of campfire kind of sound, with a lot of repetitive small impulses.

The frequency spectrum that is represented by the playback here is the light blue. And you can see that the audio is kind of dying out somewhere between, I don't know, 6 or 7 kilohertz, and then what's left is just a noise, a higher frequency noise, represented here. Now we're going to listen to the IRENE version of the same disc and what you're going to -- I believe what you're going to hear, if this would stay on the page, is obviously similar audio, less crackle, more hiss.

And that, I believe, you heard the hiss, right? So the additional hiss that you're hearing is represented in the difference, essentially, but qualitatively, between the high frequency tail here, which is also hidden under here, obviously, and the lower noise of the shellac stylus version. And, you know, that maps back to this picture that I told you, that you've got this texture, and that texture is in the material.

So let me talk a little bit about the technical aspects. So sound carriers, discs, cylinders, they're never really flat and they're never really round, and in some cases, they're way, way off that. And so this is this notion of warpage and depth of field. The 2-D and 3-D systems that directly image, they have limits, just like cameras do, to the effective depth of field that they can achieve.

For 2-D, we're limited to about 30 microns, so it's a really, really small depth of field, but having that small depth of field gets us these very high resolution images. The 3-D systems are limited to about 200 to 300 microns of depth of field. That number is typically smaller than the out of roundness and out of flatness that you often encounter in these things; even that number.

So both of these situations lead us to have active focus control systems. That's the way we deal with depth of field. We follow the surface as best we can using an additional measurement tool, let's just say, for example, a laser. Additionally, the probes that we use have sub-micron resolution, particularly the 3-D probe, this requires a very, very stable mechanics and really good vibration isolation, because it's not hard to get two things to move apart at the 100 nanometer level.

The images that we take are built up by a scanning and/or stitching method, and so when you do that, you have to have really good control of the data collection and the systematic so that everything stays in phase and things don't slip with respect to each other.

And in 2-D imaging, the illuminators that we use, they need to be DC, and I mean really DC, okay, because if they're flashing or flickering at 30 hertz, or something, that shows up as striping in the images that gets misinterpreted by the edge detection as small shifts and that shows up as spurious tones.

So those are the kind of issues that you really actually spend a lot of time dealing with once you put the thing together. So, you know, I can re-express these things as challenges. The intrinsic characteristics of mechanical sound carriers, so old, yet so precise, and one of the other speakers alluded to that, they really drive these technical requirements.

And failure to control the technical aspects can lead to artifacts in the data which are really easy to see, and, you know, we've encountered these things and we've learned how to deal with them. Vibration, AC components in the illuminator lead to spurious auto tones, poor focus control leads to crosstalk, noise, and distortion, scanning phase shifts lead to misaligned data sets, and none of these issues are fundamental in that they're impossible to solve, but they have to be handled.

And dealing with them and learning how, and they reappear when we build new systems, and they have to be worked out. That's an important part of the R&D and you shed some blood and sweat to deal with those. Okay. So the IRENE hardware and/or software has processed, already, a variety of historic materials, a number of pre, say, 1890s audio milestones, personal novelty recordings from World War II, things that soldiers sent back, aluminum and acetate transcription discs from the '30s and '40s, wax field and dictation recordings from the late 19th and early 20th century, the Edison talking dolls that the New York Times seems to be so in love with, Berliners, items from the Smithsonian Institution Volta Laboratory experimental collection, Edison tinfoils, the Schenectady foil that John McBride showed was eventually scanned, and finally, the paper phonautogram.

So I'm kind of going back in time. The present efforts and the projects as we kind of go forward include, obviously, an ongoing partnership with the Library of Congress, tools to process large collections, workflow, and diverse materials, along the lines of some of the work that I already showed, but obviously doing better at that.

Making the hardware faster and robust, and similarly, the software. We have not targeted faster than real time, but obviously, we want things to go faster. A lot of work on at-risk items, broken, and delicate materials, and how to handle them. There's an installation, recently, at the Northeast Document Conservation Center, these guys are going to tell you all about that, and then pilot projects, and even major initiatives to digitize important collections and develop best practices and specifications, and I'm going to tell you about a few of those.

So we've worked, our group, the Library, and the Smithsonian focused together on the Volta Lab collection. Some of you had the opportunity to go for the tour. If you haven't, I encourage you to go over the National Museum of American History and listen to -- or go and see Hear My Voice, Alexander Graham Bell and the development of sound recording.

In the 1880s, Bell establishes the Volta Lab here in Washington, and an instrument builder named Charles Sumner Tainter, to conduct research. They experiment with an astounding variety of materials and formats, and produce numerous patents before settling on the wax cylinder, which becomes so important. Victoria's voice gets recorded on a wax cylinder.

Right now, I think the collection is something like 400 experimental recordings and in the recent years, we had the chance to scan and process just a dozen of them, and that's what focused on the exhibit there. And I think it's been an ideal application for IRENE and a significant future opportunity to go and actually try to do this entire collection if we could ever organize that.

Carlene and Shari are the SI curators, staff folks that, I think they're going to talk to you later. I think tomorrow. I'm going to play you one record from that collection, because it's my favorite, and it's kind of the most unusual one that I've come across. One of the things that Bell did was, and this is a little off-topic, because we're talking about grooved carriers, was actually an optical recording.

He took a glass plate, coated it with emulsion, and they brought light in the mirror through a lens, and down this tube, and they exposed a light path, a light track, a spiral track, on this photodisc. And they modulated the light by spraying a jet of ink, or dye, and then causing that to move in sympathy with the sound, which varied the intensity of the light, which varied the exposure, and they wrote this exposure track.

So it starts in the middle and it spirals until it gets to this blotch here, which is an overexposed moment where, clearly, things stopped turning and the light just burned into the surface. And, you know, they're doing the usual nursery rhyme thing, and suddenly, something goes wrong, and they say something spontaneous, which I believe is a really bad word. Some people say it's not such a bad word, but I just think it's really inappropriate, so I'm going to play it.

Okay. Only in Washington. Okay. And then of course, there's a gap and they pick right back up again. And that high pitched thing was a trilled R, which was then chirped because they were speeding up. Anyway, go over to the exhibit and I think it's a great exhibit. Okay.

So a major new initiative for the project is happening now at the University of California at Berkeley, and it's called Linguistic and Ethnographic Sound Recordings from Early 20-Century California, Optical Scanning, Digitization, and Access, and it's a collaboration of the Department of Linguistics, the Phoebe Hearst Museum of Anthropology, University of California libraries, people from the Berkeley Physics Department, and the Lawrence Berkeley National Lab.

And it's a three-year project to scan a collection at the university of 2700 wax cylinders from 1900 to 1938, which were collected by Professor Alfred Kroeber and T.T. Waterman, and document

California Native American language, culture, song, and story.

The project is supported by the National Science Foundation and the National Endowment for the Humanities through a program called Documenting Endangered Languages. It's something that we're really excited about because it's going to be this very focused concerted effort that's going to run, basically, 24/7 for three years, just scanning cylinders.

And some of the technical innovations that we're trying to feature is part of this is a more efficient scanning method, so the mandrel becomes longer and we plan to stack multiple cylinders, like a shish kebab kind of arrangement, and then have a control system that can, overnight, work its way through, you know, the entire skewer.

Also, there's fixturing to hold individually broken off pieces and scan them. So I think this is going to be a really great project. Another project we're doing with Harvard is looking at aluminum discs. They were prominent in the mid-'20s to '30s as an improved material, relative to wax, for field recording, however, they have a very shallow and irregularly embossed groove.

I talked about aluminum grooves a moment ago. The groove is only 5 to 10 microns deep. The groove on a regular commercial shellacked disc is like 75 microns deep, so these are tiny grooves. These are electrically recorded, so they're typically quite loud and there's a problem to keep that needle in that groove and not fly off. So the tendency is to want to weight it down, but the aluminum is really soft.

So you're in a spot there which is not comfortable, and so optical scanning is potentially a good choice for this. And Harvard, through IMLS, has a grant for a feasibility study that we're doing with them which could eventually lead to a future initiative. This is an amazing collection from the historical point of view. It's called the Milman Perry Collection of Oral Literature.

So Milman Perry was a classicist who went to what was called the South Slavic Kingdoms at that point with this aluminum recording machine, recorded people like this individual, who were reciting poetry and stories that went on for a long time, that they learned through an oral tradition.

He analyzed the structure and the superstructure behind this that created the ability for these guys to remember these things, and that analysis of oral song led to the view that the epic poems of Homer and others were also orally transmitted because they featured the same structures.

And so this is like empirical data driven humanities work done in the 1930s looking at a living laboratory and referencing it back to these incredibly important historic -- I mean, the beginning of Western literature. So there's Perry aluminum disc and obviously, this is to evoke, I guess, it's the slaughter that occurred when Odysseus got back home and found out what the hell was going on.

And this is what one of these aluminum discs actually sounds like. They're striking when they're optically scanned. So I mean, they're pretty good; the quality of those recordings. Yes, and here's actually a display off the software package of what the aluminum data sets look like, and actually, we saw that.

It's a lateral groove and you saw the zoomed in. Interestingly, on this one, there was a little burst of noise, kind of, right in the middle of the playback, and we traced it down, and you can see that the groove is actually deformed in this one region and it corresponds to just this little patch in the material.

So lacquer transcription discs, my colleagues spoke about them earlier, they're a very, very important region, area, for optical scanning as well. There are significant collections of radial transcription discs here at the Library, and of course, elsewhere. This type of disc has a number of chemical and mechanical pathologies associated with it.

Typically, we see pretty good performance when you scan these optically relative to a stylus transfer. And a year or two back, 600 sides of these discs were done here at the Library as a pilot study using the World War II Armed Forces Radio Collection. So I'm going to first compare a stylus to an IRENE transfer on a lacquer disc from 1950. It's a radio show and then I'm going to play one of the armed forces optical transfer. So here's the stylus version.

I'm dying to know why not. Okay. And so here's the IRENE version. We still don't know. I think you could hear that the hiss -- what? Five minutes. Yes, no problem. It wasn't a question. It was an order. Okay. I think you can hear that the hiss is quite less as compared to what we heard when we listened to a shellac, so I want to remind you of the media comparison picture and the relative smoothness of those materials.

Okay. So here's the armed forces radio broadcasts. I only have the IRENE transfer for this one. So one area that is really important for lacquer and other materials are the broken discs and how do you deal with them? So the point was made that once they break, they're not in a plane anymore, much even worse than sometimes you have just with warpage, and for a direct imaging system, which has such a narrow depth of field, 30 microns, it's an issue; how do you image something that has so much, now, vertical differencing?

So here's a picture taken with the IRENE scanner of this disc, which is just broken into four pieces. Even if you break it, because the stresses change in broken things, the surfaces just move even a little bit, and with 30 microns depth of field, once you cross the crack into the next region, you're not in focus anymore.

And it's very hard for focus control to follow so quickly across such a sharp transition. So one approach that we've taken to this is to develop a special version of the scanning software that actually scans the disc in both directions. And can accommodate the crack in this way, getting one side in focus in one acquisition, and the other side in focus in another acquisition.

So with that data, we can actually tell

the control system how to end up in focus when you're here, and how to end up in focus when you're here. And now we have two images, and we have to merge them, so a process was developed, essentially, pixel-by-pixel to decide what is in focus and what isn't, and synthetically create a super focused image by merging these different images that complement each other.

Now, that works in a case like this where you're only broken into a couple of pieces, but when the number of pieces encroaches infinity, which is what often happens on lacquer discs, this doesn't work. You just cannot get things moving around fast enough.

So another approach has been under development, and one of our students is writing his thesis on this, is to take, actually, a stack of images so that, in some image, everything is in focus, you know, if you combine them, and then to use a statistical method to merge, again, pixel-by-pixel, to get a super focused image from this big stack.

So both of these use, kind of, some related concepts, but we view them as good strategies that I think are going to work out for bringing broken discs together again. See, I didn't put the bad part in. Okay. So the last point I want to make is about the fundamental technical performance of 2-D and 3-D sensors as used for direct imaging.

Sensors have spatial resolution which is well-matched to the kind of sub-micron size of groove modulations, and John talked about this, and Ottar talked about this, you know, the scale of these things. These sensors easily sample to in excess of 100 kilohertz if required, but the scan speed is proportional to sampling.

So if you want to sample at 96 kilohertz, it takes twice as long as sampling at 48 kilohertz. That's the way it works. The groove profile, okay, the amplitude, is measured on many points and the position is averaged. In 2-D, you measure a couple of edge transitions, but in 3-D, you measure a dozen, two dozen, 30 points across a groove and you can combine those.

So on a per-point basis, one of these 3-D probes, for example, confocal probe, it can measure the surface with a resolution, if you like, of something like 12 or 13 bits, just in terms of the measurement of the surface, but the measurement of the surface is not the audio. It's just the measurement of the surface.

The audio results from averaging many points and then differentiating the result, because the audio is not the surface of the record, it's the speed over which a stylus will move. So simply asking how many bits of surface height I have is not the same as what the audio is. It's, in fact, difficult to uniquely specify an audio bit depth in the same way that people do when they talk about a simple voltage converter of this synthetic result.

The real question, and I believe the focus should be on this question, is whether the measurement at the end is well-matched to the effective bit depth of the data itself. How much information is actually in this audio that we're recovering? When we're talking about historical and old materials, they're noise dominated.

There's a limit to the amount of information that's present in the vicinity of so much random noise. So what does it mean when we say that we're digitizing an audio recording to a certain bit depth? Regardless of the bit depth of the ADC, 24 bits, 16 bits, the effective bit depth, the real bit depth, is always limited by the intrinsic noise of the recording relative to the size of the signal.

So we studied, for example, recent stylus transfers that were done from field recorded wax cylinders like the Jesse Fuchs collection. When we compared the signal and the noise, we found that the available bit depth in a stylus transfer, even if it was digitized with a 24-bit ADC, is only either 7 or as little as 4 bits when you compare the noise level in those recordings to either the maximum or the average signal.

So you can use a 100-bit ADC, it doesn't exist, but when you have random noise present and a small signal, there's a limit to the amount of information that's actually there. And also, when you talk about bit depth, it also depends upon the frequency, and the distribution, and the bandwidth over which you make the measurement.

The IRENE optical scans in 2-D and 3-D are more than sufficient to measure these signal-to-noise levels as well as even lower noise levels on newer cylinders, and are, in fact, flat to 100 kilohertz if required, which I argued in this morning's presentation. The thing to keep in mind is that any measured data set has an effective number of bits, and that effective number of bits is limited by the relationship between the signal and the noise.

No matter how many bits your ADC has, the effective number of bits is limited by the random noise in the data. Okay. Finally, I want to talk a little bit about education and training. I mentioned that this morning. We have had a very significant student participation in this project. There have been about 20 University of California undergrads and about 20 thesis students from the University of Applied Sciences in Fribourg and also from Sohar and Zurich.

They have been in fields like electrical engineering, mechanical engineering, physics, anthropology, they've worked on measurements, on data analysis, on co-development algorithms, and it's been a great opportunity to expose students in engineering and physical sciences to problems in preservation and conservation.

Olivia Dill, she just graduated from Cal as physics and art history double major, and she received an undergraduate research prize that's supporting her to work on this big cylinder transfer project over this next year. So I don't know, I'm hoping that maybe she's going to go on to do preservation science with this great art history and physics background, but we'll see where it takes her. She's brilliant.

So to conclude, I hope that in the greater context, and also specifically, with IRENE, I've given you the idea that digital technology of today provides a window in this great period of sound recording that we've been focusing on.

So Bill Veillette from NEDCC, he told me, your ultimate audience is posterity. I don't know if he tells everybody that, but I like that, so I keep it in mind top. irene.ibl.gov is the project Web site and there's some -- phoebeahearstmusuemofanthropology.html has a more focused discussion of the importance of having a high resolution flat frequency distribution for capturing vertical media, so I draw you to that if you're interested. Thank you.

So Peter's running away. So we're going to do questions, right, for all of us? Okay. So I can sit down.

MR. JACOBSON: Question for Carl. This is Martin Jacobson from the U.S. National Archives. You may have said it, I may have missed it, but I'm assuming that the IRENE system uses a laser light which is close to a sine wave, is that right or are you using a frequency spectrum?

DR. HABER: I'm not sure what you mean

by a --

MR. JACOBSON: What is the wavelength of the light that you're using for reading the information off the disc?

DR. HABER: For the 2-D or the 3-D?

MR. JACOBSON: Either one.

DR. HABER: Okay. Basically, we're talking about white light. In the case of the 3-D sensor, the light source is a xenon arc lamp, which has a fairly broad frequency spectrum, and because the 3-D sensor uses this color-based confocal imaging, in fact, that's what you want. You want a lot of color so you can disperse them and then use them individually to find the light.

For the 2-D imaging, we've used a variety of white light sources also, some LED-based, some xenon-based, but some of these applications are light starved, and so, you know, having bright light illuminators that are flat in time is really important.

MR. JACOBSON: Okay. Thank you.

DR. HABER: Sure.

PARTICIPANT: We were talking about preservation and so far I don't think I've heard anyone talk about stampers. This would be the equivalent of going back to an original printed book as opposed to a photocopy book if you wanted to make a microfilm preservation copy in the old days. Why is it that, aside from the absence of stampers, and of course, many of those were melted down for the base metal, why is it that there isn't a greater degree of emphasis on going back to the ur-text, if you will, the basic text, which is the original?

And is IRENE or are the other systems that have been brought forward capable of handling stampers?

DR. HABER: Okay. There's no reason that you can't scan a stamper. It's just upside down, or depending on how you look at it. One of the items from the Volta Lab collection that is on display at the Smithsonian is actually an ur-stamper, I guess. It's like the first, I don't know, Patrick, tell me if I'm wrong, but I think it's like the first stamper, right?

It's a little laterally-cut disc that is a copper electrotype taken off wax that they must have sacrificed. So it initially worked really well. I have it if you want to hear it, but from time to time, stampers have come up. We, as you know, there are galvanos, and we did scan, I didn't have time to include it, but we scanned a galvano, which was a sacrificial galvano, so it opened up, but it's certainly possible to get a probe now that's small enough to fit inside a galvano and scan the inside surface.

So for folks that don't know, at the Berlin Phonogram Archive, the mode of operation was to bring in the new wax cylinder and immediately make a copper electrotype of it, and then sacrifice the original, and then cast duplicates for the customer or the user or the scholar, and they've recently revitalized that process, but they have 20,000 galvanos, or something, in Berlin, so that would be the cylinder version of a stamper.

But yes, there's even a little control in the 3-D analysis program for, if it's a stamper, flip it, you know?

DR. MCBRIDE: So I showed a Berliner disc, which is a sampler. They're actually easier to measure in some respects than the other discs.

DR. HABER: Right, but they're not stampers.

PARTICIPANT: But you introduced the social aspect, the university social aspect, in your slide on educational factors, and I noticed that you claimed 40 students involved with these projects. And at Berkeley, all 20 were undergraduates and no Master's or PhD students, and the thesis students were all in Fribourg.

DR. HABER: Very observant.

PARTICIPANT: Does that imply that you were not able to interest regular faculty members and departments at Berkeley in these projects?

DR. HABER: That's a great question. We've been very successful in our collaboration with our colleagues in Switzerland, and we've had many Master's and Bachelor's students. By the way, these Swiss Bachelor's students are essentially like American Master's students in terms of what they have to do. They have to write a 150-page thesis, which is original.

And we actually have had twice as many, because they also come and work with us on high-energy physics projects, so we've had a number of Master's students and Bachelor's students that all do major thesis work that take months and months.

We have never had any Berkeley professor in a technical field, you know, kind of say, let's work together or something like that. With this new project with the libraries and linguistics department, we are linked up now with professors and faculty who have a research interest in the materials.

But for example, no engineering professor has, and I think the reason why is because, you know, in major research universities, you know, people have big programs that they're already into. That's why they're there at those universities and they're typically not looking for some cool little project to, like, donate their students to.

And it's different in Europe. Europe is maybe more open-minded to this, but I think it's hard in American universities to make this kind of change.

PARTICIPANT: Two seconds. I just wanted to say that I worked in the U.K. until a few years ago, and I was actually able to get funding for PhD students to work on audio preservation, specifically, bring in doctoral students to do audio preservation projects, which, you can't do the same kind of thing here. I mean, I got an AHRC funding specifically to do that, and that's just a different way that the funding works.

DR. MCBRIDE: Well, that's humanities. In engineering, it's quite difficult. I think I did one.

DR. HABER: I mean, I've thought that there's a bunch of projects would have been great, PhD projects for students in electrical engineering, information technology, library science, it would be great to collaborate with an American colleague, but it has not happened. And, you know, I've presented this work on campus numerous times, and the statistics department, the engineering school, and so forth.

DR. JOHNSEN: I would maybe complement what Carl said. I had, in Fribourg, maybe 20 students working on their Bachelor's thesis or Master's thesis and one PhD thesis also, and in addition, yes, many of my students went to Carl. It's very easy to convince them, Berkeley, California, San Francisco, and that's why our students, some of our students, they go to Berkeley to do work for Geneva, you know, for the LHC, so I believe it's sometimes due to geography, and probably, maybe I would be more successful than Carl in trying to get Berkeley students because coming to Switzerland could be interesting for them.

PARTICIPANT: Yes, I have a question regarding, has there been any discussion about establishing a business model. In particular, if we get the universal player, the ideal thing that we have the center of excellence, and we can now go from one institution to another, and everybody is capturing the data in one way, everything you've been talking about so far is very, very heavily dependent on software. Very different than buying the needle and buying the turntable.

Once you get into the realm of software dependency, who will have ownership of the source code, and how will we preserve that and allow this, if we develop this universal tool, to continue for posterity's sake? Because if we don't have access to that, we are going to become obsolete in a very different way.

I appreciate the brilliance on stage, but like records, we all don't live forever, so how do you do that?

DR. MCBRIDE: Can I take this one? Actually, a lot of the clever software is actually in the control system for the measurement. And if you have a system, and it's doing the measurement, and you're producing the data, and then circulating the data, the data's actually quite easy to interpret. Most people with an undergraduate degree in physics or engineering could just write a piece of code, probably first year, to get the sound off.

There's nothing particularly clever in the way the sound is taken from the three-dimensional data.

PARTICIPANT: Isn't there any algorithms that are then applied? So by definition, if you're capturing raw data, level 1, think like a Photoshop model, you've got the raw data, the very next thing that happens, and you say it very casually, okay, and then we get rid of some of the noise. That's a decision someone is making and defining what particular algorithm they're going to use versus the algorithm somebody else might choose to get rid of the noise.

DR. MCBRIDE: You can download Audacity for free, and Audacity is an incredible sound system for actually taking the noise out. You can control everything. It's free on the Internet.

DR. HABER: So I agree with John that, conceptually, a lot of the transformations are not actually that complicated; however, I definitely appreciate the sentiment of your question and your comment. And I have to say that when we started down this road, because the Library of Congress engaged so early and so strongly to this, we felt like it was a good -- that was a good way to go, because here we were working with an institution that has this important leadership role, at least nationally and as a collegial institution with the British Library and the others internationally.

And frankly, mostly, we're in service to public collections. So I think that this thing really belongs in the public domain and that institutions like the Library of Congress should be the places where these tools have their ultimate home and that they're validated, and they're specified, and they're maintained, and they're distributed freely and openly, so a business model of, basically, a not-for-profit, free, open source community of users and developers guided by the Library or some international consortia of research libraries, whatever, but I think it belongs in the central, publicly-controlled, and publicly-spirited place, but yes, I think there does need to be a stability to it.

But I don't think that belongs at Apple or Google or some company that has its own agenda. It belongs in the public institution like the Library of Congress.

PARTICIPANT: First of all, I love to hear democratization of access talk. That's really great. But in general, when you start talking about archival and, you know, sharing these data sets and things like that, one of my concerns is always size of the data set. So could you provide, I know it's completely dependent upon sampling and the original source material, how quickly you have to sample along that, but could you bracket, perhaps, you know, what size of data sets we're talking about for these image sets?

DR. HABER: Yes, for example, maybe Bill or Mason, because you're producing these data sets for your customers. I mean, I can toss out some numbers, and they're going to have the gigabyte connected to them, but I don't know, is Mason and

Bill here? What are you giving out to your customers?

MR. VANDER LUGT: Cylinders are typically about a gigabyte for a 4-inch cylinder or one and a half for 6-inch.

DR. HABER: For the raw image data.

MR. VANDER LUGT: For just the raw image data. And then discs are 30 to 80, depending on the speed and size.

DR. HABER: But if you sample at twice the speed that you're doing, they're just going to grow linearly with sample speed.

MR. VANDER LUGT: That's right.

DR. HABER: So we're talking gigabyte, kind of, you know, data sets.

PARTICIPANT: When you're talking about the commercial or the workflows and you're trying to not so much commercialize, right, but send it off to production, right?

DR. HABER: Yes.

PARTICIPANT: How do you quality control? So if you're reading at three times the speed, do you still have someone that has to listen to all 15 minutes, or what do you do on that side?

DR. HABER: Okay. When you say reading at three times the speed, you're saying you're having, like, many parallel lines running at the same time?

PARTICIPANT: I guess I'm saying, once you scan it, right, somewhere you're going to scan, you're going to create the sound file, right? Who quality controls the sound file so that you can gain efficiencies?

DR. HABER: So, Peter, you pretty much designed the workflow that was used, so do you want to field that?

MR. ALYEA: As you turn your image file into an audio file, the way Carl and Earl designed the IRENE software, it effectively does what is called tracking, so you can visually see that it's lining up with the actual grooves in the media. So the biggest problem in terms of a general sense of getting an unacceptable kind of output would be when it mistracks.

Now, if you want to get into the weeds and say, you know, the absolute best you could get out of the image, that is more of a, you know, you're going to have to get in there and listen to it and tweak parameters and things. But we were trying to do it more in what would be a high throughput mode, in which case you wouldn't necessarily listen to everything, and in that case, if you get the tracking right, the audio, generally, falls into place.

And that's a fairly easy failure to detect, and so that's the way we approach it.

DR. HABER: Right, so the IRENE analysis code keeps track of how many defects, how many dust particles, how many scratches, how many conditions that fail various cut levels, threshold levels, are accumulated throughout the analysis, so you can also have a report, a histogram, a bar graph that comes out which is telling you about the data quality.

But I mean, yes, these issues, I mean, working, like, in physics, we do an experiment, we're taking data day and night, producing huge amounts, and there's, you know, quality reports that just stream out, and somebody watches and looks for outliers and things, and yes, so there's many, many approaches, and I'm sure in industrial settings, they have all kind of statistical process stuff that they do.

PARTICIPANT: I was just wondering, in the Swedish, you're very close to getting to the point of where you're going to see the results, and given what we heard, what is your, sort of, prognosis for the outcome?

MR. MOLNERYD: Yes, we have first to see the result from them in the end of August, and we have to evaluate the result that we've asked them for --

PARTICIPANT: I think the reason I asked is the second one, which was the LP?

MR. MOLNERYD: LP. That's right.

PARTICIPANT: It seemed a bit, I don't know, not exactly like I would have expected for an LP recording.

MR. MOLNERYD: They have a bit, a way, ongoing, and this is the first result they have managed, and they have, in the beginning, told us that 78 is much easier than LP, and they have started with the LPs, and when I come back here to Sweden, they are going to make a demonstration of the project so far, more precisely, and I don't know, the other two say that vinyl discs are a big problem.

And I can't say more about it, and I am very surprised from one of the companies that are very promising, and they think, oh, this will be for 78s, enough good results in the reports they will give us at the end of August. I am eager to hear. I am not expert in the -- and I take all this information with me to Sweden from Carl and all others, and mix this information together, and give the next step for us. I think so.

DR. HABER: You know, listening to the examples that you played, I was thinking that maybe they were not applying a derivative to the data that they were taking. They were just playing the shape of the groove. That's the way it sounded, and it didn't sound like they were also then applying the normal EQ curves to them, so that will definitely change the sound if they do.

MR. MOLNERYD: And what they also said is that if you have the picture, you can go on and use the picture for pre-processing, after processing the picture, and that's what I know about imaging and so on. I really don't know much more about the process to make these pictures and imaging, and I hope they have the competence to make the best result they can do with our project. We'll see.

PARTICIPANT: This is a good segue to my question, which is about, what is the bottleneck for getting faster than real time transfers? Is it the size of the light beam, or why can't it be done much faster than real time?

DR. HABER: So the 3-D probes, the best ones take nearly 200 points at once, and they take them at a certain exposure time. They just need to expose, and so they can run as fast as 2000 exposures a second for the nearly 200 points, and if you just do the arithmetic of how many exposures you have to take to cover the surface area of whatever it is, that works out to more than real time, depending upon the density.

It could be for, say, a 4-inch cylinder, could be 20 minutes or it could be a few hours, depending on what you want. So you need either faster probes, and that's going to mean they have to be brighter, basically, or probes with more points, you could imagine that. I mean, when we started, it was a 1-point probe, and then suddenly it became a 180-point probe. There's now a 190-point probe. I've heard them talking about a 1000-point probe.

But once you get to 1000 points, you start to get into issues of, you know, it becomes big and systematic things pop up. You could run multiple probes in parallel, okay, but your cost goes up by a factor of two when you start adding. Probes are like $50,000.

For the IRENE and VisualAudio 2-D scanners, one of the bottlenecks is just like disc-writing speed for the data you're generating. So that's -- the 2-D scanning, I think, has the greatest, because right now, it takes like 15 minutes to scan a 3-minute disc. I think there's hope along those lines of pushing that, but maybe, Earl, you want to comment because you've thought about this a lot. Where do you see it going?

MR. CORNELL: So we're definitely, sort of, a factor of two speed limited by disc writing currently, and that's just because we haven't taken the time to buy or research faster. And then of course, we're still light limited as well. I mean, once you get past that factor too, then it becomes a light limiting factor.

DR. MCBRIDE: So can I make a comment here? I mean, the only reason for increasing the speed is throughput, but if you automate, then you don't need it. You know, if you automate the system, and you can just have the disc filed off and then you can just feed them through, so actually what's more important is actually automation, so you just automate the whole thing.

MR. MOLNERYD: Our input in the beginning was that we have -- will take more than 20 years for us to make these discs we have in the archive for audio file, and we have it continuing for five years if they put the records on and it takes five years, and we have --- what can we do? The material in the archive will be standing in the shelves and no one has listened to it.

Is it possible to find a new method to make this and listen to them? And this was the input for us to see, is it possible, and they had to prove it, the companies who's left in the project, and so on. And one input is also, is the quality enough just for listen? We know that vinyls and 78 records are stable. You can take it back if you don't crush it.

And you have a listening copy for researchers and so on, not in the high-end quality, but it's enough, perhaps, for --- hear what they say, and so on. his is one of the results that we have to manage in the end of it.

DR. HABER: So if you just want, I guess you would call them access copies, for listening purposes of 78s, you could take a scanner like the VisualAudio scanner or the IRENE scanner, which are essentially the same thing, you could put four or six cameras on there and just --- you could just scale it up and you could get it in. If you could deal with the disc writing, or you have multiple disc drives, and you just dump the thing.

I mean, that's like a brute force. If your goal is really to read a three-minute record in a minute, you probably could make that happen by just throwing multiples of $6000 at it until you had enough cameras. I mean, at some point you won't fit any more cameras. It becomes a packing problem.

PARTICIPANT: Question for Dr. McBride, in your Journal of Acoustic Society paper, you note that you are getting 10 nanometers of axial resolution on a top to bottom groove dispersion of 10 microns, which is 1000 units of measurement, which is about 10 bits, right? Simple math. And you state that the digital audio bit depth is equivalent to the image bit depth. Do you still hold to that?

DR. MCBRIDE: Have you got the paper there?

PARTICIPANT: I do.

DR. MCBRIDE: Okay. Well, when was that published?

PARTICIPANT: 2008.

DR. MCBRIDE: Okay. Wow.

PARTICIPANT: The question is, as advanced as -- is there -- I'm understanding from Dr. Haber that the axial resolution in the image domain does not translate to digital audio resolution. I'm trying to resolve that.

DR. MCBRIDE: You need to take the derivative of the surface to get the sample, but you still need to have the raw data to start with.

PARTICIPANT: Of course. No, I've only taken a couple college level classes in physics, but --

DR. MCBRIDE: You still have to have the raw data.

PARTICIPANT: You have to have the raw data there, that's right, and if the raw data is only there with a precision of, in your case, 10 bits.

DR. MCBRIDE: I don't recall. On that

slide I showed you, it was the 20 kilohertz harmonic, oh no sorry, 2 kilohertz harmonic in the spoken word.

PARTICIPANT: Yes. If the image is measured with the precision of 10 bits, is there any transformation that you can do to increase resolution?

DR. MCBRIDE: You cannot go --- you cannot change the resolution of the sensor that you're using. So that graph was actually about defining -- it's about the quality of the potential audio that you can get from this, and we were interested in preservation, so we're at the first stage of this. It's about preserving the raw data for future use.

PARTICIPANT: Understood.

DR. MCBRIDE: And that's where that measurement came from. It was a displacement measurement against frequency. If you look at the lower frequency, I think I pointed that out on the slide, it was 600 hertz? And the magnitude of that particular harmonic was much higher, yes? So this is where I was focusing on the harmonics associated with the spoken word, yes? So it's a very kind of detailed scientific issue.

PARTICIPANT: I understand. I do understand. But let me ask the question again, perhaps in a way that we can all understand. If something's measured with 10 bits of precision in the digital optical domain. That's a level of precision. Is there any transformation that can be done, because that's the digital --- that's the preservation master, right? Is there any transformation that can be done to bring that up to 16 bits.

DR. MCBRIDE: Okay, hang on. The 10 bit, we didn't publish that. You've done a calculation, which is the ratio of the 50-nanometer height, which we've just referred to.

PARTICIPANT: No, you published it. I'm looking at it right here.

DR. MCBRIDE: Does it say 10 bit?

PARTICIPANT: That's my point of contention.

DR. MCBRIDE: I don't think that's

correct, because it actually -- is that comparing the depth of the groove to the 50 nanometers? Is that the number that you said?

PARTICIPANT: Yes. I'll read it to you so we can have it.

DR. MCBRIDE: Yes, I don't think that's correct.

DR. HABER: Okay. Here's the thing, you measure the surface at a certain point, but you take many points and you combine those points to get an estimate of the depth of the groove at a certain moment in time. So the depth of the groove at that moment in time is synthesized from a number of measurements. So by combining multiple measurements through an averaging process, you can improve, through averaging, your measurement.

So just taking the single --- just interpreting what the resolution is on a single point and then comparing it to some number, 10 microns, 20 microns, that you're sort of picking out from someplace, is just the beginning of a more complicated story that leads you to a certain performance which comes out from combining all this data, taking the surface, differentiating it, and turning it into audio.

A statement that you can make, which is very well defined, about the instrument that you buy is that it has a range of a certain, say, 200 microns, 300 microns, depending on the instrument that you choose, the optics that you choose, it has a certain resolution inside that range, 50 nanometers, 10 nanometers, 120 nanometers, depending on the one you use. And that ratio, between its fundamental resolution and its range, that defines its capability in bits --

PARTICIPANT: Theoretical bit depth.

DR. HABER: -- for a single point measurement. Okay?

PARTICIPANT: Right.

DR. HABER: That's just the beginning of a more complicated story.

PARTICIPANT: I understand.

DR. HABER: And there's no -- sorry. It's very hard to see you, because Mason's blocking, it's a little disconcerting. I feel like I'm talking to somebody behind somebody, but you know, so that's why I went into this discussion because there's no simple little answer, and frankly, when people in audio engineering say, oh, I've digitized this at 24 bits, that's, it's a big simplification about what's actually going on.

PARTICIPANT: I understand. Unfortunately, this conference --

DR. HABER: Just a minute, there has to be a recognition that noise, random noise, is a --- sets the scale that we have to base everything on. If you want to measure the radius of a billiard ball, I can give you a pair of precision calipers and you can measure that radius to whatever you like. If you want to measure the radius of a tennis ball, the diameter of a tennis ball, it's a poorly defined problem because a tennis ball is furry, alright?

So there's an arbitrariness due to the random structure of the surface of a tennis ball. That means that, fundamentally, you can't measure it the same way you can measure a billiard ball.

Audio engineers recognize this when they say, well, a 24-bit ADC is really only 20 bits, or 21 bits, because, you know, there's fundamental noise in the ADC electronics that knocks, already, 4 or 5 bits out of the game.

And I tried to build up that explanation in just a few minutes. It could be the subject of an hour-long discussion. You know, you can go in Wikipedia and read pages and pages about -- you know, you can -- if you want to run your clock at 3 megahertz, you could have a 1-bit ADC and digitize audio at very, very high quality.

So there's no simple little -- I know you're trying to get a simple little answer and say, well, I can do 24 bits and you can only do 7, or something like that, but it doesn't work.

PARTICIPANT: Carl, that's not what I'm trying to do.

DR. HABER: Okay.

PARTICIPANT: I'm trying to get to the bottom of a question that's been bothering me and a lot of other people, and for which you and Dr.

McBride seem to have very different answers.

DR. MCBRIDE: Okay. I wasn't, I said --

DR. HABER: Do we?

DR. MCBRIDE: Well, I don't think we do. I mean, maybe a line that's in a paper which was published, what, eight years ago?

PARTICIPANT: Okay. I understand. Okay. That's fine.

DR. MCBRIDE: You should have given me a prior warning. I'd have looked up the background for it. The sensor -- I just did a calculation, the sensor we use is actually 15-bit resolution. You've got a 350-micron range and it resolves to 10 nanometers.

PARTICIPANT: Yes, and that's where Carl got his 12 to 13 bits as well.

DR. MCBRIDE: I think what we were saying is that, it's really important that if you want to get the best quality audio from these old recordings, you have to have a sensor that has a minimum resolution of around, is it, 50 nanometers? What does it say? You've got it there.

PARTICIPANT: Yes.

DR. HABER: 50 nanometers, whatever it is.

DR. MCBRIDE: Yes, 50 nanometers.

DR. HABER: And that's something in that part. It depends on the material, but that's sort of the --

DR. MCBRIDE: That's the key thing.

MR. ALYEA: We're cutting into the next session and it seems like maybe we should deal with this, I definitely want to deal with this, but maybe John could look at the paper again.

DR. MCBRIDE: Yes, sure.

MR. ALYEA: We could pick this up. We have a lot more time. At the end of thing we have more time. We're cutting into other people's talks right now, so I think we should have a very quick coffee break, about five minutes, and move the new talkers in and then --- sorry, David.

(Whereupon, the foregoing matter went off the record at 3:48 p.m. and went back on the record at 3:59 p.m.)

MR. ALYEA: Okay. So our first presentation is, we actually have kind of, two duos, this time. From NEH we have Joshua Sternfeld and Jesse Johnston, so Joshua is speaking first, right? Yes.

DR. STERNFELD: That's the best kind of introduction. Jump right in. So, hello, everyone. I am Joshua Sternfeld. I'm a senior program officer with the National Endowment for the Humanities, Division of Preservation and Access, and I'm here with my colleague, Jesse Johnston, also from the Division of Preservation and Access.

So today, we are going to talk about a few things here. I think we're going to move from the very technical discussions that we've heard up to this point and perhaps in the course of just talking about our grant programming and some of the projects that we supported over the years, we will begin to kind of see a way of answering some of the questions that Carl was raising earlier this morning, about where do we go from here, and how do we take all of the research and technologies that have been developed up to this point, and begin to scale them upward and begin to apply them to more and more collections.

So that will, for our purposes here today, that will begin with a very, very brief overview of NEH. Many of you are already familiar, some of you are actually awardees, past awardees, but I realize that there are some that are from outside the U.S., so I will give you the 30-second rundown of what we do. And that will kind of move us into some of the projects that we have supported along audiovisual-related projects, mostly through research and development, but as Jesse will indicate, we have had support through almost all of our grant programs.

And so one of the things that I think will become more and more clear as we talk about those programs is that, perhaps we could look at the suite of programs that we offer as a kind of paradigm for how to begin to support these new approaches to preservation and access going forward.

And Jesse will conclude with a little bit of indication of how interested NEH is currently with the topic of audiovisual preservation and access. We are ramping up our interest more and more as we speak and so we're delighted to be here and we thank the Library of Congress for having us, and hopefully we will be able to continue this dialogue both in this panel and perhaps afterwards.

So a little bit about NEH. We're divided into a collection of offices and divisions that covers the scope of humanities activities, and that includes everything from your very basic individual fellowships, and research, and educational opportunities, to development of public programs, things like your documentaries, museum exhibitions, and so forth, and what Jesse and I work on, the Division of Preservation and Access.

So we actually are devoted to helping to preserve humanities-related collections of all types: audiovisual, archaeological, literary, historical, et cetera. And so in the Division of Preservation and Access we try to cover the entire range of activities as best we can, that is, related to the basic care and stewardship of humanities collections, for all types of institutions, I should indicate.

I know there's some, perhaps, misperception that we're only geared towards the high research universities, but that's simply not true. We do our best to also reach out to smaller institutions, things like historical societies, public libraries, small museums, et cetera.

So that range of activities, of course, can cover everything from your basic intellectual control, things like cataloguing, arranging, describing; things like your basic conservation treatment, reformatting, of course, rehousing of collections.

We do digitize. We have actually gone into the academic world and heard that, well, NEH doesn't digitize any longer. That's not true. We love to digitize audiovisual collections and see projects of that sort. We also do support the development of databases, spatial tools, reference resources, and we're increasingly, in the last, I would say, five or six years, moving towards sustainable solutions for preventive conservation care.

And that --- those are projects to address the entirety of a storage environment. And we actually have a few, a handful, of projects related to audiovisual collections on that front as well.

So a little bit about where we come from in terms of our support for audiovisual preservation and access. And, mostly, my job here today is to focus on the R&D side of things, since I think that's kind of most in tune with what we're talking about here today.

So we can go back all the way to the 1980s and you probably won't find a program specifically called R&D like you do today, but in various forms over the years, we supported very kind of specialized research projects. At the beginning of the '80s and '90s, a lot of those awards went to the Image Permanence Institute over at the Rochester

Institute of Technology. Many of you are probably familiar with the great work that they do.

They began in the mid to late '80s working on -- actually, early '80s, working on photographic prints, and then it's almost like a chronology of audiovisual formats and the projects that they've worked on over the years, moving through various formats, landing in magnetic media in the '90s, and then in the early 2000s, we started to see support for a number of other institutions.

Many of you may be familiar with the Sound Directions project that was a collaboration between Indiana University and Harvard, and those were a series of projects that eventually developed a set of audio digitization guidelines, so we're very proud of that. One of the reasons we're proud of that project is not only for the R&D work that they did, but the fact that it had such an impact on the wider cultural heritage field.

And so what we love to see are projects in our other grant programs, particularly, for example, in projects to digitize specific collections that are referencing the Sound Directions guidelines as the standard that, or best practice, that they're trying to adhere to. So when we see R&D work show up in footnotes in our other grant programs, it gives us a little chill, or thrill, if you will.

And also around this time, in 2004, we got this kind of curious project called Developing an Optical Method to Recover Sound from Mechanical Recording Media. I was actually able to dig the original file up and it's quite a massive collection of documents.

You see here one of the original images for what then became IRENE. So NEH was proud to be at the kind of ground level for the work that Carl and his team did beginning back in 2004. And obviously, you can see the fruits of all that labor here today. So I won't go into the history of the project since we've already gone, you know, quite extensively, but I thought, you know, seeing the early 2000s kinds of technologies is kind of a fun little trip down memory lane.

But we do lunge the IRENE project as a kind of exemplar of research and development projects that we've supported. One is simply that it responded to a real need in both preservation as well as cultural heritage fields, and I think all of the speakers that you've heard this morning covered those needs quite well.

And then I think a true testament to the success of an R&D project is that it then, of course, leads to additional projects, additional research, and that could be in, specifically, an R&D-type setting, or it could lead to, as it did with additional testing and evaluation, that could also lead to reconsiderations of broader stewardship practices, and as Carl was hinting at, shifts in how we think about education and training as well.

So that leads me a little bit to talking about a few of our main grant programs coming out of the Division of Preservation and Access. And again, for some of you, this might be of interest directly in thinking about getting a project to us. And of course, whenever we go out, we always love to field questions, one-on-one types of questions of a project that you have in mind, but I think also for the purposes of this panel and for today, again, the suite of programs that we have can also be considered a kind of framework for considering how to expand all of these great technologies going forward.

So the way we -- this is one way of how we look at our grant programming. The core of what we do, the foundational program that we have, is called Humanities Collections and Reference Resources. And that's a program that is geared to enhance the preservation and/or access of specific humanities collections.

So we get, again, institutions of all shapes and sizes come to us and say, we have a great collection of X and we need to do work, whether it's to rehouse them, to catalogue them, to reformat them, et cetera. And for us, that really is our wheelhouse, our humanities collections, because it enables us to bring in humanities practitioners, scholars, curators, folks in the cultural heritage sector, into our panel and review setting, and conversing with preservationists of all stripes, archivists, librarians, and so forth.

But we realize that the work that goes into preserving a humanities collection requires a lot of other work that kind of encircles that work. And I'll be talking, again, a little bit about the research and development and the particularities of that program, sustaining cultural heritage collections, as Jesse will explain, tackles the preventative issues as well, and of course, there's education and training, and we're very proud to support a few audiovisual-related projects in that realm as well.

So, R&D. We underwent an intensive review of this program in the last year and a half or so. We just received --- and we restructured the program and we just received that first batch of applications a few weeks ago, so that's beginning the review process. So in one sense, the good news is that the next deadline is an annual deadline, and you have until June 21st of 2016 to think about getting another project to us.

Again, this program is designed to address major challenges in preserving and providing access to cultural heritage collections, and it's also designed to support the development of standards, best practices, methodologies, and workflows. It supports applied research. And I know some of the subsequent presentations coming up will be talking about considerations for how to develop best practices in this area.

So we have two tiers of funding, and I'll give you the quick specifics around those, but I do want to add, we are very proud to support international collaboration. The only requirement being that the primary applicant institution needs to be a U.S.-based institution, but we have supported and will continue to support projects that show collaboration with institutions outside the U.S., so again, something to consider going forward.

So I had mentioned that we have two tiers of projects here. So a Tier I project is for up to $75,000, and what we notice -- this is really the newest aspect of our R&D program -- what we notice is that if we only had one tier, which was up to $350,000, that there can be a lot of great work at these lower levels that are getting squeezed out from the limited funding that NEH has.

These can be anything from surveys, to case studies, to small-scale experiments. So particularly in the field that we're talking about here today, optical image preservation, I think there could be a lot of cases where these kind of small case studies could really develop valuable data that could be fed back into the larger community.

And one of the things that we're trying to emphasize with this level as well, is that this can be a great opportunity for graduate and post-doctoral training and experience.

Secondly, we have a Tier-2 level. Again, this is the higher level for up to 350,000 for up to three years, and this is the whole enchilada, if you will. We want to see projects that are really striving to develop best practices standards, methodologies, workflows, that can reach a very broad preservation and cultural heritage audience.

So one of the things that we changed with this program, and this might be something that we could pick up again during the Q&A and beyond, is this notion of dissemination. So we have a lot of great projects coming out of R&D over the years and we always did require a short section that we called dissemination, but over the last few years we were getting dissemination plans that included, well, we'll develop a blog that we will kind of update every once in a while, and then we'll get on Twitter and we'll tweet it every now and then, or on Facebook.

And somehow, this whole community is going to just magically pick up all this great new knowledge and run with it. And in some cases, that was true, and it worked, but I think in a lot of cases the dissemination plan failed to match the great work that was being conducted at the R&D level.

And if you are going to develop standards and best practices, you have to really engage the community, right? You have to really understand what the community needs are and how they're going to receive that knowledge, and work with it, and begin to accept those best practices going forward.

So we tried to elevate our dissemination plan a little bit. We require a one to two-page appendix that -- and we're asking a kind of demonstration that you understand how to reach your community that you define within your application.

And of course, there's a whole suite of possible dissemination activities that can fall into a kind of creative plan, which can run from workshops and training, to educational models, publications, then we get into things like code sprints, and conferences, and then of course if you're looking at getting a standard accepted, of course, that will require communication with a kind of regulatory body in some cases.

So the bottom line is, we would love to see R&D projects that are interdisciplinarian. And again, I think this came up in some of the earlier discussions about ways in which the work being conducted at a very technical level can link up with some of the work of the cultural heritage field.

And I think the project that Carl mentioned earlier, the linguistics project, which I think Jesse will talk about further, is a great example. So we're striving to see projects that are interdisciplinary in nature. International, again, that sort of merge the humanities and the sciences. And of course, the sciences, when we're talking about preservation, can run the gamut from the information sciences, computer science, natural sciences, of course, in many cases here today, and conservation science as well.

So I feel --- I think I would be remiss if I didn't actually give an example of an optical audio project, besides supporting IRENE in its early days. And so we've actually supported a project coming out of the University of South Carolina, which they call AEO-Light. Is anyone familiar with AEO-Light? A little bit, yes. Okay.

So this is, again, why dissemination is so important. This is actually related to the film community, so I don't know if I'm committing heresy talking about film in an audio media environment, but I think it's applicable here. So as you can see in this image here of a still from a film of Martin Luther King, Jr., you see the optical track running on the side.

So what the project director, who is Greg Wilsbacher, he is collaborating with the mathematical institute at South Carolina. What they're doing is essentially very similar to what we've been talking about here today. They are creating images, digitized images, of 60mm, 35mm film, which captures both the optical sound as well as the image. They're even going so far as to go all the way out to the sprockets, and I'll show you why in a second.

And then what they do is run an algorithm that extracts the optical sound and links it with the digitized film images. Apologies for my email going off like that. So why is this significant? Why was it funded by NEH? Well, quite simply, what they're hoping to accomplish, and I think they're getting to it by their second phase of their project, is essentially cutting, perhaps not completely in half, but close to half, the time it takes to digitize a film.

The process, traditionally, is to capture the image first and then audio second, and then link them up afterwards. They're saying by capturing audio and image simultaneously, and then, linking them through this mathematical software program, you can digitize the film in half the time. That's the hope, and I think they're getting along their way.

So if you go to their site, you can actually download their software program. It's open source. I think they would be thrilled to hear from a lot of you to get your feedback on that. And of course, I'll play just a very short clip of a test clip that they have on their website here, so bear with me. Hopefully this will work.

And this is a film of Benito Mussolini from 1931. I believe from their Fox Movietone collection.

(Video played.)

DR. STERNFELD: So you can see the audio track on the left-hand side running simultaneously with the image, and you can see, of course, the words on the very far edges, and the argument goes that, from a kind of digital humanities standpoint, this is valuable information. This is, essentially, providential information that certain humanity scholars, media scholars, film scholars, might find quite valuable to help date and locate the origins of their film.

So with that, I'm going to turn off my email and hand it over to Jesse.

MR. JOHNSTON: We rehearsed that. As Josh said, he's talking more about the research and development sorts of projects that we've supported that are somewhat related to the IRENE and audio imaging efforts, and I'm going to talk about the rest of our programs and hopefully give you an idea of some of the ways that they can support some of the efforts that are happening in the audio preservation world.

The Humanities Collections and Reference Resources Program, which is one that Josh has already mentioned, is our largest program, and it's really focused on advancing preservation and access for specific collections, as Josh said, or, also, the design of reference resources, like dictionaries and grammars, and in both of these areas, it's actually supported audio projects.

So as you can see, we're sort of illustrating what these programs can do by way of example, but if you have specific questions about what sorts of activities are eligible or ineligible, technically speaking, we can answer those questions at the end. The examples, I think, are probably most interesting, so I'll focus on those now.

In 2012, the University of North Carolina received an award, many of you are probably aware of the Southern Folklife Collection there, and they're preserving audio recordings and photographs that document traditional music and musicians of Southeastern United States.

And that actually is a hybrid

collection, so there's a lot of audio there, but this is an archival collection that also has photographs and related field notes as well. And so hybrid collections are certainly something that we frequently see, given that our collections really cross the broad range of documentation of humanities activities.

In the reference resources category of this program, we've also supported discographic projects, including what's become the American Discography Project, which grew out of the Encyclopedic Discography of Victor Recordings at the University of California in Santa Barbara, which received an award in 2011 and a couple of other times.

Recently, we've supported a new level in this program called Foundations, which offers support for, sort of, early-stage projects. So if all of the technical details, perhaps, aren't set, hoping to bring together, maybe, technical preservation people with humanities scholars to lay the foundations, as it were, for future projects.

I also want to tell you about the education and training program that we offer. This is a program that also, like R&D has been a longstanding priority at NEH and we've supported preservation education and training in various forms since the 1980s when we took over a program of regional field services that was previously funded by the National Endowment for the Arts.

And these now offer basic preservation services, such as preservation assessments, rehousing supplies, conducting emergency preparedness workshops, and you're going to hear a lot more from one of those services in a moment. And they can sort of be the example.

But this program also supports other things. It supports Master's Degree programs and it also supports workshops that address preservation and access topics of national significance and broad impact. And we include there in our guidelines that the preservation of and provision of access to recorded sound and moving image collections is one of the priorities in this program.

I think that it would be accurate to say that we hope to continue that priority, and also, to give you an idea, we have supported that in the past. For example, we've supported an award to New York University in 2014 that provides scholarships to students in the Moving Image Archiving and Preservation Program for doing up to 24 internships, that is 24 students doing internships, at cultural heritage institutions in the New York Metropolitan area.

The Center for the Conservation of Art and Historical Artifacts also has offered audio workshops through their preservation field service, as have some of the other ones.

The Documenting Endangered Languages Program is a unique program that we have. It's a partnership with the National Science Foundation. Carl has already mentioned the award that we recently made in partnership with the National Science Foundation to the University of California, Berkeley. And this is, I guess, at least the second award that NEH has made which supports an IRENE-focused project.

And as you know, they're working on digitizing cylinders of Native American languages gathered in California. It's a very unique collection with high value to the humanities for documenting early history of anthropology as well as, and this is why it came to the Documenting Endangered Languages Program, documentation of languages that are no longer spoken or they may be in a sort of, what linguists call, moribund state, and very few speakers, or an aging speaker population.

We feel that these are invaluable for scholarly research and broader purposes of cultural and linguistic revitalization. And as you heard when the project was mentioned earlier, this has actually supported an internal institutional collaboration between different units of the University as well as Lawrence Berkeley Labs. And so we hope that is an example that points to sort of the power of these programs to help to encourage that sort of cross-institutional and inter-institutional work.

I'll just say a few words in conclusion about our priority on audiovisual preservation. You have a better sense now, I think, of our division's main grant programs and how they relate to preservation and access of audiovisual materials. And we have placed a priority on the preservation of audiovisual content through various efforts, some of which I've covered here, since at least 2000, so for about 15 years.

And in 2001, I believe, we supported a symposium here at the Library in cooperation with some of the other agencies in the Library on preserving folklore documentation. And that was kind of a first step and we hope that we're continuing in that vein.

Our hope in sharing with you some of the audiovisual projects that we've discussed today, and that NEH has supported over the years, is partly to remind you that our grant programming can support a wide spectrum of preservation activities, from research and development, to education, to digital reformatting.

We also want to communicate to you that we are interested in exploring options for participating and supporting this effort in better ways. We have heard the need to consider areas as far ranging as format and playback equipment obsolescence, mass collection assessments, selection and appraisal, cold storage for films, professional training, and continued research and development.

But we do hope to hear from those working in these fields directly about how we could better support that effort in a continuing way in the future.

However, in a restricted resource environment, we do hope that in the spirit of the National Recording Preservation Plan, and other calls to action that we're aware of as funders, we do want to facilitate the maximum impact of what is a small amount of financial resources by fostering collaboration and coordination among many stakeholder communities, which is why we've stressed that theme of collaboration throughout.

Many people that we've spoken to have reminded us that audiovisual collections are often overlooked, whether that's for lack of funding, preservation experience, or simply lack of interest in relation to other types of collections, or perhaps lack of technical familiarity among collection stewards.

Given NEH's unique connection to the scholarly and professional communities, we think that one of our contributions as a funder can be to bring together the voices of scholars and researchers, and other audiences for these materials, with preservation communities. That's sort of one of the unique positions that the Division of Preservation and Access occupies.

In fact, we believe that only with the voice of content and subject specialists, including many that are here and who may be speaking tomorrow, can effective arguments be made for the significance and importance for expending scarce resources that are required to select, assess, reformat, and use the many audiovisual resources that face uncertain futures.

And finally, by reaching out to various user audiences, including researchers and educators, we hope to raise greater awareness of the scale of the challenge of audiovisual preservation, which, as many of you know, is great. And we also hope to foster clearer understandings of the various uses, approaches, and ways to make choices that confront us as we face this audiovisual preservation challenge.

If you're interested in our programs, you can ask us questions after this. You can also contact us on email, and here's our contact information, so thank you very much.

MR. VEILLETTE: Hi everybody. I'm Bill Veillette. I'm the Executive Director of the Northeast Document Conservation Center. Very quickly, we're, in some respects, the odd man in the room because we're actually a paper and book conservation center. And you might ask why are we involved with audio preservation?

Well, back in the '70s, we started doing microfilming and film duplication. And then in 2008, because it was hard to get film, and I'm very impressed that you can still get it in Switzerland, we transitioned the imaging over from microfilming and film duplication to digital imaging. And then at that point, because there's so much involved in post-processing with a digital transfer, it didn't become a very big leap for us to start thinking about audio.

And in fact, most of our clients are archivists, not, obviously, audio engineers. And many, many archivists are also charged with the stewardship of audio collections, but they're scared to death. So here we are as a non-profit that's been serving them for over 40 years. They're very comfortable with us and trust us, and they started asking us to get involved with audio, so that's why we are here today.

That led to a conversation with the Institute for Museum and Library Services to take a look at this thing called IRENE, this is a few years ago, and we ended up writing a grant to do a pilot to see if we could take IRENE from the lab, successfully, to the market, so to speak, and use it not as a scientific instrument, but as a reformatting production tool.

So these were the principal goals of the project, there were other questions as well, but I thought those other ones wouldn't be as interesting to this audience. But you can read them here. We were curious about these kinds of questions because these are the questions that go into whether we could sustain a service and run it without losing money and going out of -- going bankrupt.

So I'll just cycle us through these questions. First of all, what we did find was that it is actually quite easy to learn to operate the IRENE system with a person of average intelligence, however you want to define that. It starts to become a little bit more difficult when there are mechanical issues, so you do need some greater-than-normal mechanical aptitude to do some troubleshooting. If you are all thumbs and have no mechanical aptitude, this is probably not the technology for you at this stage in its development.

But in terms of the software, although software is very stable, it's not particularly buggy, which we felt very good about. We were concerned that it might be crashing every ten minutes and we can't get anything done. That is not the case at all. But when you want to make an enhancement, or a tweak, or something, that requires a call back to Lawrence Berkeley National Labs.

So almost half-jokingly, but not really, we have said that we need to take "Key Man" life insurance out on Earl Cornell, which is a little bit of a morbid thought, Earl, sorry about that, but we haven't done it. Don't worry. Okay. In terms of the hardware, we did receive this upgraded model, you've all heard about that, you saw pictures of it that Carl presented, it has the 2-D and the 3-D on the same platform so that we could switch easily to 3-D to scan a disc without having to go through a lot of extra setup.

There are some issues though. As parts

get upgraded or replaced, there are some issues with having to then potentially have to modify the software and even the firmware, the OEM might have to make some modifications, so we saw that with, I think, our probe, you know, early on, where we had to go back to the OEM and ask them to make some changes.

But the good news is that it's very reliable. Once it works, it works, mechanically and in terms of the software. Now, there are some challenges and we'll segue to Mason in a second here where he'll talk about some challenges that had more to do with the carriers themselves and having to do some fussing because of the optical approach.

Okay. What I really want to focus on here is some of the questions that were being asked early in the day, which is, so, if you're trying to use this as a tool, how practical is it? How much can you get done? And in our case, our clients wanted to know, what is this going to cost?

Now, there are three variables to providing this kind of a service. One is your efficiency, which is, how many billable hours out of all the paid hours in the day your staff can work. The next is the billing rate, which has more to do with NEDCC's overhead. And then finally, the productivity, which is, ultimately, the ultimate measure, which is how many recordings you could do per billable hour.

So, in terms of efficiency, it actually has nothing to do with IRENE or any piece of equipment that you're using. This is how many hours out of the day you're able to bill your time out. So you're taking out any administrative time and you're reducing it down to, how many hours am I actually serving a client?

So that really has to do with the operator, you know, their work ethic in the workplace. We start work at 8:30, are people serving clients at 8:31 or are they kind of getting around to it after a couple cups of coffee, right?

The billing rate also, nothing to do with IRENE, has all to do with the costs of labor and overhead, but productivity has to do with two things. It has to do with the operator and their workflow habits within the constraints of the tool.

So what are the constraints of the tool? Well, in our case, one of the things we were trying to test with the pilot was, you know, could we operate multiple IRENE machines at once by one operator? And our conclusion at this stage is, with our current configuration of the machine, it would be impractical to try to have more than one machine per person, and you'll see why in a second. I have a slide.

But, we need to be mindful that in the, you know, 12 or 15 years of development of IRENE, the focus has not been on its production capacity. It's been on the science of imaging and analyzing those images. So there's this now new phase of opportunity to start focusing, in addition, on productivity. So these are not “problems” with IRENE, these are “opportunities” going forward.

There is a little bit of multitasking you can do. Once you've scanned your first carrier, you can be scanning the second carrier while you're analyzing the first one. Okay. So there's a little bit of 2:1 workflow for all but the first and last recordings. And then because of that, the first and the last recordings are 50 percent more expensive.

So all this means is that we say to clients, look, if you have more than one carrier, don't send them to us one at a time. Bunch them up and help us help you by making these projects as -- you know, scaling these projects up a little bit because we can get you --- save you some money.

So, you know, here's what we can do with our one machine. We can start a scan and then while "Recording 1" is being analyzed, we can start scanning "Recording 2," et cetera, until finally, at the end of the day, you know, you're just kind of analyzing the last one. Now, you could be recording something for the next morning, of course, but this is just an example of if you had an eight carrier project.

Now, 2:1 with IRENE is not the same as 2:1 with a stylus. 2:1 in real time means that you can't be doing two things at once, or rather, you're attending to two things at once, so one of those things that you're not paying attention to is not being paid attention to.

But 2:1 with a scanning approach means that the imaging can be unattended, which is okay, because there's really nothing to do while it's being imaged. I mean, you could watch it as it's being imaged, but there's no point in doing that. But the analysis is always attended. So there's not this issue that people are concerned about, you know, 12:1 workflows because the vendor's really not paying attention to what they're doing, right? They're just more focused on their own bottom-line.

So here's an example. At NEDCC, we strive for 90 percent efficiency. In other words, we want 90 percent of the seven and a half hour workday to be focused on clients and not on admin stuff and things that would be distracting. So that translates to 6.75 hours a day. Our billing rate's $125 an hour, and if it takes 90 minutes to scan and analyze a single carrier, using that graph that I showed you previously, we could do eight carriers a day.

So if you do the math on that, it's $105 per recording. Now, it can get quite a bit more expensive if the recordings are broken, or cracked, or dirty, we have to clean them, you know, there could be added costs if they're large discs, bigger cylinders, and those kinds of things.

So there is this opportunity, as I mentioned, to, now start to --- now that we know the technology works, and in some cases it works better than the stylus approach because we can get more signal, this is this opportunity now to start focusing, in addition, on productivity. And you've heard Carl talk -- he gave you the picture of the shish kebab design, which we're very interested in, but he also just verbally described some other approaches you could do by just throwing cameras at the situation.

I was talking with Earl about this at lunch and, you know, we were talking about some of the challenges and the bottlenecks, and what might seem expensive in terms of equipment from our standpoint is a cheap date, because when you're billing out at $125 an hour, you know, if you had to spend $500 for some component, or $1000 for some component, to cut the time in half, that's nothing.

You know, you can make that back within hours, you know, or days. So really, this isn't a physicist's math, it's a business person's math, but I would not discount spending a lot of extra money, frankly, to speed up the process because you'll get it back in production. We just need to do the back of the envelope math.

So at this point, what I'm going to do is turn it over to Mason who's going to describe some of the challenges with the physical carriers that can throw a big-time wrench into that graph I showed you where we could do eight units per day.

MR. VANDER LUGT: Hi, everybody. Can you hear me? I'm Mason Vander Lugt and I operate the IRENE at NEDCC. I'm going to try make this quick because I've got a presentation of my own and we're kind of going over. So none of us expected IRENE to be ready off the shelf when we got it. We understood we were kind of inheriting complicated technology that was combining image sensors, lasers, vibration isolation, networking, so we expected a lot of that and had a lot of small problems that were mostly kind of bug fixes, hours on the phone with Earl and Carl.

But kind of one of the most standing issues that we'd still like to tackle is making lacquer disc scanning a little bit more reliable. So the biggest problem with lacquer discs is that they can be difficult to image, and we saw this earlier in Stefano's presentation, that the groove bottom isn't always visible and that's where we want to get our audio from.

So this is actually a shellac disc on the corner of the screen, and I'll explain a little bit more about how these images work in my own talk, but what you see is a number of grooves, the vertical black lines, in the middle of each black line is another white line, which represents the bottom of the groove.

So, you know, when I see that I know I can get good sound from these images. Lacquer discs are cut instead of pressed, and so in the middle of the groove instead of a clear groove bottom, we often see nothing. If it's an angle, it's doesn't really reflect light, so we see this, you know, just a black line.

So this is an example of a lacquer disc. One of the ways we've tried to get around this problem is by using side lighting. So instead of illuminating the imaging from the top, we bring in a light from the side at an angle, and it gets a different part of the groove, and these kind of thicker white lines are the bottom corner of a groove, I guess.

It can act as a groove bottom and we can get good audio from that. The problem is that it takes a lot longer. Another way to get around this problem is to analyze the groove tops instead. Again, this takes a lot longer and doesn't produce the same kind of audio we expect from a groove bottom. Let me go back a minute.

So these are things that we'd like to see become more reliable, using either side lighting or the groove top analysis, we can usually get something from a disc. It's just not the eight recordings a day we like to keep productive with. Another problem we have with lacquer discs is a loss of focus. Lacquer discs, the lacquer itself is transparent or translucent instead of the kind of opaque and textured shellac discs.

So the lasers we use to detect the recording's surface and move the camera to stay in focus sometimes see through the lacquer and read the bottom of the disc or see texture in plasticizer oxidation. And because IRENE uses such a small depth of focus, even very slight misreadings in the key-in system, which is the laser, can cause an instant loss of focus.

And there are a couple ways we've talked about working around this, like increasing the depth of focus or using a different system to detect the recording's surface. And then finally, one of the ways we lose time frequently with lacquer discs is a difference in tracking them.

As Carl described, tracking is telling

the computer software where the groove goes through the image, so from, kind of, top to bottom of the image and across. Lacquer discs can be a huge range of different types of recordings, really. So from commercial masters that are predictable and well recorded to, basically, toys and paper-based voice-o-graphs.

And I'll get into this more in my next talk as well, but this is a shellac disc, and it's a commercial disc, but often, we see something more like this where I don't know if the recordist wasn't trained or the machine wasn't calibrated, but the grooves run together.

And then when the software is trying to track this automatically, it doesn't have the same information and kind of skips grooves or, you know, reads in-between the lines and creates these kind of blue things, and then I have to go through and manually track. So this is the Fourier tracking algorithm and the other automatic tracking we use is new tracking and that uses the groove bottom, and so often, if there isn't this groove bottom or if there's texture in the top of the groove, like we see in this scratch here, it just skips grooves.

And so again, I can work around this using manual tracking, it's just one of the things that cuts into our time if we're trying to keep any kind of meaningful workflow. So I'll turn it back into Bill and I'll be back up in a minute.

MR. VEILLETTE: So, you know, another consequence of some of these issues with the carrier that you don't know are there until you've scanned it, is that, you know, one thing clients expect is when you give them a proposal they want to know that the price is the price. So, you know, we do have some caveats in the proposal to say, you know, sometimes we can't see the groove bottom, et cetera, and so far, so good.

You know, people have been pretty understanding that the technology is, you know, still being developed. Okay. Just a couple slides to end. We got a big grant from the Mellon Foundation to do a study of the audio preservation market because we're planning on getting further into audio preservation by reformatting magnetic media.

And so we hired AVPreserve to do this study and this is just a little data point out of it where they found that there are 46 million preservation-worthy grooved carriers that require a specialized workflow. Now, this 46 million happens to be the same number as the 46 million that the Heritage Health Index came up with for all audio back in 2005.

AVPreserve's number now is 250 million audio items. And there's a paper where they can defend that number that I'll just refer you to, you could Google it. But, you know, this 46 million isn't all rare grooved carriers like wax cylinders and lacquer discs. But presuming we just use that number, if we can do 1,720 cylinders or discs per year based on 215 productive days in a best case scenario, and you do the math over 50 years, you would need 535 IRENE machines over 50 years to complete all this work, if you did them all with IRENE machines, which is not appropriate.

They're appropriate for certain

materials and less so for others. But the point here is that prioritization and selection is going to be critical. So our counsel to clients is that IRENE is best used for delicate recordings that are unique, rare, recordings that are rare formats where the equipment may not be available to play them otherwise, and then the obvious one that everybody thinks of and we don't need to tell them, is damaged and broken media.

But we specifically say to them, please don't, you know, use it for commercial recordings. We want to reserve our time on our IRENE machine for people who have these other types of materials.

The final thing that we were charged with in our IMLS grant was to figure out if there was a way we could help sustain R&D for IRENE. And there were lots of different things we were talking about, but this turns out to be, really, probably the best way we could help.

We had some focus groups that were not part of the IMLS project or part of that Mellon-funded study, but we kind of folded some

IRENE questions into it and found out that clients are willing to pay up to 15 percent more for an optical approach. And remember, most of these clients are archivists, they're not audio engineers, and, you know, these are people who don't like anybody to touch anything.

So it's very appealing to them, the non-contact approach, and they're willing to pay a premium for it. The question is, can we use this 15 percent to support R&D and would that provide enough money to move the needle. So if you do the math on our billable hours, and our rate, and whatnot, the bottom-line is it comes out to $27,000 per IRENE machine that we could feed back to the effort to develop the system further.

So that might seem like a small amount, but really, the world needs more than one IRENE machine providing services, and so it could end up being a pretty good amount as you multiply this by the number of machines that are put into production. Finally, our study found out that 46 percent of prospective clients could not proceed with their audio projects unless they had a grant.

So if you do the math on that as well, knowing that based on the previous slide we need $208,000 per year to sustain each IRENE machine at full capacity, we need to make sure that there are enough grant programs out there to the tune of $96,000 per machine to sustain the machine, otherwise, we actually get starved of, kind of, capital, which is the oxygen for doing any kind of work.

So that's our presentation. If you want to learn more, we have a very robust section of our website where we describe the service in more detail, what the deliverables are, you know, what is part of the proposal, these caveats that I mentioned, and we also have some very good examples about how to clean discs that we put up there based on LC's vetted approach and how to ship, and pack, and all that kind of stuff.

So it's a very interesting website in terms of, you know, understanding how an optical scanning service is being presented to prospective clients, so I encourage you to go there to look at those kinds of things.

So at this point, I'm going to turn it over to Mason, who has a really interesting presentation. I know we're running late on time, but let me just pull us here and get his up and we'll get him going.

MR. VANDER LUGT: All right. So I want to talk a little bit about using images for not their intended purpose. What am I going to talk about? So we use IRENE because it is contact free and doesn't risk our recordings, but as an intermediary step, we are left with these digital images that are very high magnification and high resolution, and I think these images are pretty useful if we take a step back.

I think they can teach us about the recordings themselves, the degradation of the recordings, and also help us improve the operation of our IRENE system. So I think, some of you --- everybody understands this at this point. IRENE uses a 2-D and a 3-D camera. They're both line scan cameras and so we take a pass, the cylinder rotates, and we move on to the second pass, and over time, we're left with a complete map of the surface of the recording.

So the natural proportions of the images, their width is constrained by the image sensors, but we want to take enough image samples so we can get the desired audio sample rate. So this long black line on the left is actually the natural proportion of the image. If we zoom in at full resolution, we're left with this middle image, and then there's another kind of image, I'll show it today, that is a little bit less easy to understand, but we call them binned images, and this is what we do when we're analyzing the software to help us -- analyzing the data rather, to help us track.

And we sample the images vertically so we can get a full picture of where the grooves go without having to scroll through a long image like this. So this is just a condensed, vertically condensed image like the one on the left.

So what can we learn from the pictures?

One of the things is objective information, like groove width. So this is a series of four different types of disc recordings, organized in descending order by their groove width. The top is a Voice-o-Graph disc and the bottom is a microgroove disc, so they range from 110 micrometers at the top to 50 micrometers at the bottom.

We can also gain qualitative information, so when I see these, I know what kind of audio we'll get from the recording. On the left, it's a commercially produced wax recording, and on the right, it's a field recorded wax recording, and just based on the training of the recordist, the conditions in the field where it was recorded, and the type of machine it was recorded on, it changes the audio quality we get.

And when I see these images, I kind of get an idea of what to expect from the audio. So like I said before, this is a binned image of the groove, so you can see where the grooves go throughout the image. This is a commercial shellac recording, so this is as about consistent as it gets.

The grooves are evenly spaced and they all have groove bottoms, I believe.

This is the image that I used in my other presentation. Sometimes they are not recorded as well and the grooves move throughout the image. This is another example. This is from a paper-based lacquer disc and the groove spacing is terrible, it wasn't planned well for the excursion of the grooves, and I don't know what you would get if you played this on a turntable, but if you look at it under IRENE, you know that you're not going to get something that you want.

This is another kind of poorly recorded disc. You can see the cutting lathe was not set to record deeply enough, so once a revolution it actually lifted from the disc, leaving a section where it was undercut. This is something, again, if you played it on a turntable, you would probably come to understand that something is wrong with the disc, but we see very clearly now what's going on.

And, you know, beyond assessment, one thing we can do with IRENE is actually just trace across these gaps and at least get all of the information that is recorded. Obviously, we can't get anything that wasn't.

Here's another example. I turned the image on its side, just because I think it's interesting. It was kind of undercut or maybe the recording machine wasn't working at its best that day, but the head actually skipped across the surface. Each of these black lines is a groove, and at some point in the recording process, it skipped across the surface like a stone.

This is a cylinder that was overcut and this is another one of these we're comparing with a mechanical transfer. We wouldn't have known exactly why this happened playing it with a stylus, just that, in two sections, it would get quiet. But looking at it under IRENE, we can see exactly why it got quiet, it was because it was overcutting and the stylus wouldn't have read this fully.

This is just a fun example. Can't really learn much from it, but it's interesting to look at, I think. And another tinfoil where you can see, kind of, the nature of the impression on the foil.

So my examples up until now have been about what we can learn from the original recordings as they were recorded, but we can also learn a lot of interesting things about how recordings get damaged or degrade over time.

This is another paper-based lacquer disc. This is the first couple of grooves, kind of a run-in groove, a lead-in groove, and you can see that it's just a mess. It's in tatters. It wasn't recorded well, but you can also see, through subsequent playbacks, I guess high tracking, poor styli had dragged across it and created extra grooves.

I think this actually was recorded twice. I think this is another cutting head that was put on the disc and cut a new track. As we progress through this, we see more new tracks, more wear, and as we progress, we see some kind of chattering up in the top, some inconsistency.

If we zoom into that, we can see the disc was actually bent. The paper is more flexible than the lacquer, and so when you bend the paper-based lacquer disc, the surface lacquer crazes and leaves us with something that we couldn't get a whole lot of audio from. Here's another example of a crazing disc.

You can see the actual grooves in this image don't look like the grooves in most of our IRENE images, and that's because it's actually a white disc. That leaves us with less contrast and that's another challenge for us is figuring out how to image these and get the same contrast that'll allow us to get good audio.

So continuing on in the damage and degradation, this is a disc that was delaminating, and you can see, it's a lacquer disc that was delaminating. I think you all understand what that means, but the surface material, the lacquer and the bases don't always stick together well over time, and the pieces kind of can separate from the surface and separate from each other.

And zooming into this, you can see that the places where it began delaminating, the lacquer was actually pulling apart and it separated at the thinnest part, which is the groove bottom. I think that's interesting. So we can learn something about the nature of the delamination process.

Here's another image of a flake of a lacquer disc that was delaminating. There's a bit of an illusion here as the camera loses focus below this line, but you can see this piece was flaking up and didn't keep in focus, and you can see the light shines through the translucent lacquer.

Here's another condition I think a lot of you are familiar with. This disc was exuding plasticizer and looking at it under the magnification of IRENE, we can actually see the crystalline structure of the fatty acids that came out, and it's interesting to me to see that the same structure is present beneath the groove bottom, so that's throughout the disc.

Here's a broken glass-based lacquer disc where you can see -- so the disc broke, but also chipped away at the interface of the break, and you can see the shear of the glass chip. And also, if you look closely, you can see the texture of the glass, the bubbles in the glass.

Here's a frayed paper-based lacquer disc. This is a Voice-o-Graph disc. It was a voice letter from the second World War that had been played too many times and wore away the lacquer in this area, and you can see the paper base beneath it, but you can also see the lacquer was kind of pulling away, and the thickest part in the groove tops remained.

You can kind of see this when you look at the disc at eye level, but this gives us a lot more information about how this happened. Here's a disc that was covered in adhesive residue and left glue covering the groove bottoms and the groove tops. We couldn't see through this at all. And a similar one where a disc had been written in with wax crayon and it filled the groove, but left the groove tops alone.

And here is a disc that was damaged with some kind of mysterious pitting. I don't really know what caused this, but the disc was very noisy. I think you can actually see the disc is splitting again down the middle of this groove as well.

Groove top wear. This is what it looks like when we play a disc and it's noisy. And so I'll go through a couple of similar examples for cylinders. This is what we expect to see from a clean commercial cylinder. They're vertically recorded, as you all know, so the discs, the grooves don't go side-to-side as in a laterally cut recording, but up and down, and that's why you see the pits in here.

This is what mold damage looks like under IRENE. You see the audio grooves on the left side of the recording, on the left side of the image, and the mold damage on the right side. One thing I think is interesting about this is that the mold seems to have taken up residence in the pits created by the audio.

This is, you can see impressions made by the fibers of original wax cylinder recording container and this is actually useful information for us. It tells us that you don't really want to store wax cylinders in the original cotton-lined containers because they can leave marks on your disc that will be audible later.

Similarly, this is a brown wax cylinder that was beginning to effloresce, or bloom. This is kind of -- I don't think there's a lot of research on this, but this at least shows me that this is not mold. It's actually a different failure mode. Here's a talking doll cylinder recording, as Carl mentioned earlier. This is what they look like, kind of, typically, and then you can see one that had been played more.

I don't know if the first has been played never and the second once, or the first only a few times, and the second many times, but you can see the wear in the middle. It's these black lines where it had been played by a stylus.

So finally, we can use the images and also images produced by the analysis software to assess our IRENE performance. The reason I have this image here is because this motor, if it's off in alignment by fractions of a degree vertically or in the Y-axis, into and out of this image, we can lose focus throughout the image.

And one of the ways we can calibrate it so that doesn't happen is by using the depth information gathered by the 3-D probe and the same information in another plane gathered by the KEYENCE laser. So in order to prevent artifacts in our sound in 3-D, we want our images to align very precisely. It sounds like it would just be 360 degrees around and the width of the sensor in a step, but that leaves us with a small duplication in the data.

So what we do is guess and check until there's no overlap, but the features of the recording are visible smoothly across the lines at the end of a pass as you see on the left side, in-between passes as you see on the right side. So this is what we want to see from tracking. When I began working with IRENE in fall of 2013, we couldn't leave the software to analyze on its own because we didn't have any indication whether the tracking went successfully, so we asked them to add tracking images so we could leave the software to analyze overnight, and review tracking in the morning.

This means I can do other things instead of watching the software analyze. This is what we have at risk if we don't know whether the tracking was successful, but if I see this looking at tracking images in review, I can go back and redo this and it won't give us faulty audio.

We can also see how blob clean is doing. I don't know if anybody has described blob clean at this point, but it's a way of cleaning up the sound only, because the software can tell, based on the shape of the pits, in 3-D, and the depth, what's mold and what's audio. If we are too aggressive with this algorithm, it can remove audio, but we can review the images while we're analyzing and see that this one actually did a pretty good job figuring out what was blobs, or mold spots, and what was audio, and so we know we're using the right settings for the blob clean algorithm.

So this is a dirty lacquer disc, not really exuding plasticizer, mostly just dirty, but we can't really get great audio from this, so when

I began working with IRENE, we looked into ways we could improve our data. Fortunately, somebody from preservation research and testing here was already working on this, and we adopted a cleaning method from his process, and ended up being really successful with it.

This is the same image, or the same disc, rather. You can see the features here. The only difference is before and after cleaning. So this tells us that our cleaning is doing well. If you want to know more about our cleaning process, I've got a URL here, but you can just go to our website and it's easy to find.

And finally, like I said earlier, we have added side lighting to our process for discs without a good groove bottom. When we originally invented this, or designed this, rig to hold the light in its place to illuminate the side of the groove wall, we overloaded the motor and it left us with an image that was out of focus when the camera had to move quickly.

We figured out that we could get rid of

some other weight in the system and found immediately that that helped our focus problems. And I think, finally, we can see in our data, before we analyze it, if there are bugs in the system and don't waste time analyzing the data, so when I see these glitches on the right side in the detailed view, I know that I have to fix something before I try to keep analyzing.

So thank you. If you have any questions, I'll be here. And --- I'll be here.

PARTICIPANT: About what percentage request 3-D analysis? Do you have an update on the --

MR. VEILLETTE: Well, the clients don't --- can you hear me? Yes. The clients don't actually request 2-D or 3-D. We use the approach that's best for that carrier, but we've had only one client that we're proposing to use 3-D on for their discs, but it's a vertically cut disc too, so.

MR. VANDER LUGT: We do it sometimes with paper-based lacquers and we did it once with an aluminum transcription disc, but yes, as Bill said, we don't let people choose which process they're going to use, we just use whatever makes sense.

PARTICIPANT: Well, no, my question really went to how many people want all the information that's in a 3-D analysis? Because this morning, I inferred that if you really want to preserve a document, you want to preserve all the data you can, and that made the 3-D analysis desirable, but what you're telling me is that your clients are not cognizant of that.

MR. VEILLETTE: Right. Well, there are image files for 2-D and 3-D. And the 2-D image files are also very interesting to look at, as we just saw. We have two different kinds of templates for clients. We have one for institutions and one for private clients. We've had, actually, quite a large number of private clients initially because there was more publicity through popular media.

But now the mix is shifting toward institutional. With the institutional clients we, as a default, give them the image files and charge them for the external hard drive, you know, that can hold all those files. With the private clients, we don't, as a default, but we let them know that we're not giving them the image files, but if they want them, we could add that to the proposal. Most private clients don't care about the image files, they just want the audio.

But we also have language in our proposal that tells the institution to treat those image files as the master file and to take care of it. We also let them know that there is no software that they can use to do anything with it, but we will make our best efforts to maintain our software and that the ultimate goal is to do what Carl was talking about earlier, to have somebody, perhaps the Library of Congress, maintain an open source version of it.

DR. NYE: I'm Jim Nye from the University of Chicago and it was wonderful to hear that at NEH, international engagements are taking an important place in your planning. I'd like to press that a little bit further, particularly because I think, arguably, the linkages that NEH has had with JISC, DFG over the years have been enormously productive and also have extended the funds that are available in limited quantity from Congress.

And my question is about what might be involved to initiate new relationships with counterpart bodies, let's say, in the developing world? And I work in South Asia, so I'm thinking specifically, what would be involved in setting up a relationship like the DFG/NEH, but with India?

DR. STERNFELD: Yes. That's a good question. It's a long and tenuous process to do a kind of longer term relationship like the one you're describing. I don't know the history of the DFG program and how long it took. I do know that we are undergoing some talks. We're having a transatlantic, it's not a conference, a coordination amongst, I believe, 12 funding bodies from both, I think, South America and Europe, as well, is where you get the transatlantic moniker.

And so I think that's in the early stages. I do know that our director of the Office of Digital Humanities actually went to Korea recently to, sort of, probe possibilities of creating some Asian connections. I don't think I can speak much further than that. It's all sort of, you know, we're just sort of probing possibilities at this point.

The program that has perhaps the most international in its structure is the Digging into Data Program, which has had two funding cycles so far. That's a coordination of, I believe, eight funding institutions, some U.S.-based, NEH, IMLS, NSF, some of our Canadian counterparts, some of our U.K. counterparts, and so forth.

And, yes I know, for example, Carl, I guess he walked out, but he was proposing, for example, a kind of pilot study of the possibilities in a kind of digital humanities sense of what we can do with these data sets, these large data sets. That Digging Into Data Program is, I think, a perfect vehicle to encourage collaborations, international collaborations, to take these data sets and to derive new ways of understanding them and analyzing them.

MR. JOHNSTON: Well, and also, if I could just add, I think the implication of the question is that all of these examples aren't reaching some areas of the world, I think. And typically, those are initiated with particular people, you know, and so if those agencies in those areas are interested, there are probably people that they can contact at NEH, if you want to talk further.

(Whereupon, the meeting in the above-entitled matter was concluded at 5:21 p.m.)